



**MVS/370
Linkage Editor and Loader
User's Guide**

Program
Product





MVS/370
Linkage Editor and Loader
User's Guide

Data Facility Product 5665-295
Release 1.1

GC26-4061-1

Second Edition (December 1985)

This is a major revision of, and makes obsolete, GC26-4061-0.

This edition applies to Release 1.1 of MVS/370 Data Facility Product, Program Product 5665-295, and to any subsequent releases until otherwise indicated in new editions or technical newsletters.

The changes for this edition are summarized under "Summary of Amendments" following the preface. Specific changes are indicated by a vertical bar to the left of the change. These bars will be deleted at any subsequent republication of the page affected. Editorial changes that have no technical significance are not noted.

Changes are made periodically to this publication; before using this publication in connection with the operation of IBM systems, consult the latest IBM System/370, 30xx, and 4300 Processors Bibliography, GC20-0001, for the editions that are applicable and current.

References in this publication to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM program product in this publication is not intended to state or imply that only IBM's program product may be used. Any functionally equivalent program may be used instead.

Requests for IBM publications should be made to your IBM representative or to the IBM branch office serving your locality. If you request publications from the address given below, your order will be delayed because publications are not stocked there.

A form for readers' comments is provided at the back of this publication. If the form has been removed, comments may be addressed to IBM Corporation, P.O. Box 50020, Programming Publishing, San Jose, California, U.S.A. 95150. IBM may use or distribute whatever information you supply in any way it believes appropriate without incurring any obligation to you.

PREFACE

This publication supports MVS/370 Data Facility Product and provides application programmers with the information necessary to use the linkage editor and loader to prepare the output of a language translator for execution. Additional information on the operation and use of the linkage editor and loader is directed to the system programmer responsible for installing and maintaining the operating system.

ORGANIZATION

This publication contains an introduction and two major parts:

- "Introduction" defines the functions and gives recommendations for the use of the linkage editor and loader.
- "Part I. Linkage Editor" describes the processing facilities and operation of the linkage editor:
 - Chapter 1, "Overview," describes object and load modules and gives a general overview of linkage editor processing.
 - Chapter 2, "Uses of the Linkage Editor," provides descriptions of the functions of the linkage editor, and explains its relationship to the operating system.
 - Chapter 3, "Defining Input to the Linkage Editor," describes how to define the primary input data set, how to use the automatic library call mechanism, and how to include other data sets as input.
 - Chapter 4, "Specifying JCL to Run a Linkage Editor Job," explains the job control language necessary to run a linkage editor job step.
 - Chapter 5, "Specifying an Operation with Control Statements," summarizes the various linkage editor control statements that can be used in running the job.
 - Chapter 6, "Editing a Control Section," describes how to change external symbols, replace control sections, delete control sections or entry names, order control sections or named common areas, and align control sections or named common areas on page boundaries.
 - Chapter 7, "Invoking the Linkage Editor," gives the macro instructions used by a problem program to invoke the linkage editor.
 - Chapter 8, "Interpreting Linkage Editor Output," describes how to interpret the output load modules and diagnostic information produced by the linkage editor.
 - Appendix A, "Sample Linkage Editor Program," contains four sample programs illustrating the use of the linkage editor.
 - Appendix B, "Linkage Editor Requirements and Capacities," describes the record-processing capacities of the linkage editor, the types of devices that can be used for the intermediate data set, and the amount of virtual storage the linkage editor requires.

- Appendix C, "Designing and Specifying Overlay Programs," describes how to use the overlay facilities of the linkage editor to minimize the amount of virtual storage required.
- "Part II. Loader" includes function descriptions and operating instructions for the loader program:
 - Chapter 9, "Overview and Uses of the Loader," describes the functional characteristics of the loader, its compatibility with the linkage editor, and the restrictions on its use.
 - Chapter 10, "Preparing Input for the Loader," explains how to prepare an input deck for the loader, including how to specify EXEC statements and how to use DD statements to define loaded program data.
 - Chapter 11, "Invoking the Loader," shows how to use the EXEC statement or specified macro instructions to invoke the loader program.
 - Chapter 12, "Interpreting Loader Output," describes how to interpret the diagnostic and error messages, and the optional storage map, produced by the loader program.
 - Appendix D, "Loader Storage Considerations" describes the virtual storage space required by the loader program.
 - Appendix E, "Loader Return Codes," lists the return codes that can result from loader processing and defines their meanings.
 - Glossary
 - Index

The diagnostic messages issued by both the linkage editor and the loader program are described in MVS/370 Message Library: System Messages, Volumes 1 and 2, GC28-1374 and GC28-1375. The description of each message includes an explanation, a system action, and a problem determination action to be taken.

PREREQUISITE KNOWLEDGE

In order to use this book efficiently, you should be familiar with OS/VS2 MVS job control language.

REQUIRED PUBLICATIONS

You should be familiar with the information presented in the following publications:

- MVS JCL, GC28-1300, describes the job control language used to run the linkage editor and loader programs.
- MVS/370 Message Library: System Messages, GC28-1374 and GC28-1375, describes the diagnostic messages issued by the linkage editor and loader programs.

RELATED PUBLICATIONS

Within the text, references are made to the publications listed in the table below:

Short Title (as it appears in the text)	Publication Title	Order Number
Data Administration Guide	<u>MVS/370 Data Administration Guide</u>	GC26-4058
Data Areas	<u>OS/VS2 Data Areas</u>	SYB8-0606
Initialization and Tuning Guide	<u>OS/VS2 MVS System Programming Library: Initialization and Tuning Guide</u>	GC28-1029
Installation: System Generation	<u>MVS/370 Installation: System Generation</u>	GC26-4056
JCL	<u>MVS JCL</u>	GC28-1300
Linkage Editor Logic	<u>MVS/370 Linkage Editor Logic</u>	LY26-3921
Loader Logic	<u>MVS/370 Loader Logic</u>	LY26-3922
Routing and Descriptor Codes	<u>OS/VS Message Library: VS2 Routing and Descriptor Codes</u>	GC38-1102
Service Aids	<u>OS/VS2 MVS System Programming Library: Service Aids</u>	GC28-0674
SMP System Programmer's Guide	<u>OS/VS System Modification Program (SMP) System Programmer's Guide</u>	GC28-0673
Supervisor Services and Macro Instructions	<u>OS/VS2 MVS Supervisor Services and Macro Instructions</u>	GC28-0683
System Codes	<u>OS/VS Message Library: VS2 System Codes</u>	GC38-1008
System Messages	<u>MVS/370 Message Library: System Messages, Volumes 1 and 2</u>	GC28-1374 GC28-1375
TSO Command Language Reference	<u>OS/VS2 TSO Command Language Reference</u>	GC28-0646
Utilities	<u>MVS/370 Data Administration: Utilities</u>	GC26-4065

NOTATIONAL CONVENTIONS

A uniform system of notation describes the format of linkage editor and loader control statements. This notation is not part of the language; it simply provides a basis for describing the structure of the commands. The command format illustrations in this book use the following conventions:

- Brackets [] indicate an optional parameter.
- Braces { } indicate a choice of entry; unless a default is indicated, you must choose one of the entries.
- Items separated by a vertical bar (|) represent alternative items. No more than one of these items may be selected.
- An ellipsis (...) indicates that multiple entries of the type immediately preceding the ellipsis are allowed.
- Other punctuation (such as parentheses, commas, and spaces) must be entered as shown. A space is indicated by a blank.
- **BOLDFACE** type indicates the exact characters to be entered, except as described in the first four bullets. Such items must be entered exactly as illustrated.
- Lowercase underscored type specifies fields to be supplied by the user.
- **BOLDFACE UNDERSCORED** type indicates a default option. If the parameter is omitted, the underscored value is assumed.

SUMMARY OF AMENDMENTS

RELEASE 1.1 LIBRARY UPDATE, DECEMBER 1985

SERVICE CHANGES

Rather than reflecting system capability, the structure of this publication reflects the customer's approach to a specific task.

The title of this publication has been changed from MVS/370 Linkage Editor and Loader to MVS/370 Linkage Editor and Loader User's Guide.

All other MVS/370 titles referred to in this publication have been changed to their corresponding MVS/XA titles.

Information has been added to reflect any service changes.

RELEASE 1.1 UPDATE, MARCH 1984

NEW DEVICE SUPPORT

4248 Printer

Information to support the IBM 4248 Printer has been added to the DATAMGT and IODEVICE macros, and to the description of SYS1.IMAGELIB. Device tables and the example of sysgen in Appendix A have also been changed to reflect the 4248.

3262 Model 5 Printer

Information to support the IBM 3262 Model 5 Printer has been added to the DATAMGT and IODEVICE macros, and to the description of SYS1.IMAGELIB. The 3262 Model 5 printer is specified in the IODEVICE macro as UNIT=4248.

RELEASE 1.1, OCTOBER 1983

NEW DEVICE SUPPORT

3800 Printing subsystem Model 3 Compatibility

The IBM 3800 Printing Subsystem Model 3 is now supported in both compatibility and full function mode. Information to support the 3800 Model 3 has been added to the DATAMGT and IODEVICE macros. Device tables have also been changed to reflect the 3800 Model 3.

Five new system data sets, SYS1.FDEFLIB, SYS1.FONTLIB, SYS1.OVERLIB, SYS1.PDEFLIB, and SYS1.PSEGLIB, are also documented for the 3800 Model 3 with the Print Services Facility (PSF).

3430 Magnetic Tape Subsystem

Information to support the IBM 3430 Magnetic Tape Subsystem has been added to the IODEVICE macro. Device tables and the example of sysgen in Appendix A have also been changed to reflect the 3430.

4245 Printer

Information to support the IBM 4245 Printer has been added to the DATAMGT and IODEVICE macros, and to the description of SYS1.IMAGELIB. Device tables and the example of sysgen in Appendix A have also been changed to reflect the 4245.

NEW PROGRAMMING SUPPORT

Graphic Access Method/System Product (GAM/SP) Release 2

Information to support GAM/SP and the 5080 Graphic System has been added to the DATAMGT and IODEVICE macros. Device tables have also been changed to reflect the 5080.

System Modification Program Extended (SMP/E) Release 1

Information has been added throughout this manual to support using SMP/E in place of SMP.

SERVICE CHANGES

Additional editorial changes have been made throughout the manual, and the preface has been reorganized.

NEW DEVICE SUPPORT

The device tables in Chapter 4 and Appendix B reflect support for the IBM 3380 Direct Access Storage Models AE4, BE4, AD4 and BD4.

CONTENTS

Introduction 1

Part I. Linkage Editor 3

Chapter 1. Overview 4

Object and Load Modules 6

 External Symbol Dictionary 7

 Text 8

 Relocation Dictionary 9

 End Indication 9

Linkage Editor Processing 9

 Input and Output Sources 9

 Load Module Creation 10

 Assigning Addresses 10

 Resolving External References 11

Chapter 2. Uses of the Linkage Editor 12

Linkage Editor Input 12

 Links Modules 13

 Edits Modules 13

 Aligns Control Sections or Common Areas on Page Boundaries 14

 Accepts Additional Input Sources 14

 Reserves Storage 15

 Processes Pseudoregisters 15

 Creates Overlay Programs 16

 Creates Multiple Load Modules 16

 Provides Special Processing and Diagnostic Output Options 16

 Assigns Load Module Attributes 16

 Allocates User-Specified Virtual Storage Areas 16

 Stores System Status Index Information 17

 Traces Processing History 17

 Lengthens Control Sections or Named Common Sections 17

 Assigns an Authorization Code to Output Load Modules 17

 Assigns Addressing Mode 18

 Assigns Residence Mode 19

 AMODE/RMODE Combinations from the ESD 20

 AMODE/RMODE Hierarchy 21

 Assigns Read-only Attribute 21

Relationship to the Operating System 21

 Time Sharing Option (TSO) 22

Chapter 3. Defining Input to the Linkage Editor 23

Primary Input Data Set 23

 Object Modules 23

 From Cards 24

 As a Member of a Partitioned Data Set 24

 Passed from a Previous Job Step 25

 Created in a Separate Job 26

 Control Statements 26

 Object Modules and Control Statements 27

 Control Statements in the Input Stream 27

 Control Statements in a Separate Data Set 28

Automatic Library Call 28

 SYSLIB DD Statement 29

 System Call Library 29

 Private Call Libraries 30

 Concatenation of Call Libraries 30

 Library Control Statement 30

 Additional Call Libraries 31

 Restricted No-Call Function 31

 Never-Call Function 32

 NCAL Option 32

Included Data Sets 33

 Including Sequential Data Sets 34

 Including Library Members 35

 Including Concatenated Data Sets 35

Chapter 4. Specifying JCL to Run a Linkage Editor Job	37
EXEC Statement—Introduction	37
EXEC Statement—Job Step Options	37
Module Attributes	38
Downward Compatible Attribute	38
Scatter Format Attribute	38
Not Editable Attribute	39
Only-Loadable Attribute	39
Overlay Attribute	39
Reusability Attributes	40
Refreshable Attribute	41
Test Attribute	41
Authorization Code	41
Addressing Mode Attribute	42
Residence Mode Attribute	43
AMODE/RMODE Combinations in the PARM Field	43
Default Attributes	43
Incompatible Attributes	44
Special Processing Options	44
Exclusive Call Option	44
Let Execute Option	44
No Automatic Library-Call Option	45
Space Allocation Options	45
SIZE Option	45
DCBS Option	50
Output Options	51
Control Statement Listing Option	51
Module Map Option	51
Cross Reference Table Option	52
Alternate Output (SYSTEM) Option	52
Incompatible Job Step Options	52
EXEC Statement—Region Parameter	54
EXEC Statement—Return Code	54
DD Statements	55
Linkage Editor DD Statements	56
SYSLIN DD Statement	56
SYSLIB DD Statement	57
SYSUT1 DD Statement	57
SYSPRINT DD Statement	57
SYSLMOD DD Statement	58
SYSTEM DD Statement	59
Additional DD Statements	60
Size Parameter Guidelines	61
Cataloged Procedures	62
Linkage Editor Cataloged Procedures	62
Procedure LKED	62
Procedure LKEDG	64
Overriding Cataloged Procedures	65
Overriding the EXEC Statement	65
Overriding DD Statements	66
Adding DD Statements	66
Chapter 5. Specifying an Operation with Control Statements	67
General Format	67
Format Conventions	67
Placement Information	68
ALIAS Statement	69
CHANGE Statement	70
ENTRY Statement	72
EXPAND Statement	73
IDENTIFY Statement	74
INCLUDE Statement	76
INSERT Statement	77
LIBRARY Statement	79
MODE Statement	81
NAME Statement	83
ORDER Statement	84
OVERLAY Statement	86
PAGE Statement	88
REPLACE Statement	90
SETCODE Statement	92
SETSSI Statement	93

Chapter 6. Editing a Control Section	94
Editing Conventions	94
Changing External Symbols	96
Replacing Control Sections	97
Automatic Replacement	98
Example 1	98
Example 2	99
REPLACE Statement	101
Deleting a Control Section or Entry Name	102
Ordering Control Sections or Named Common Areas	103
Aligning Control Sections or Named Common Areas on Page Boundaries	105
Chapter 7. Invoking the Linkage Editor	107
Chapter 8. Interpreting Linkage Editor Output	109
Output Load Module	109
Output Module Library	109
Member Name	110
Alias Names	111
Entry Point	112
Authorization Code	113
Residence and Addressing Modes	113
Reserving Storage in the Output Load Module	113
Processing Pseudoregisters	114
Multiple Load Module Processing	114
Diagnostic Output	115
Diagnostic Messages	115
Module Disposition Messages	115
Error/Warning Messages	116
Sample Diagnostic Output	117
Optional Output	117
Control Statement Listing	118
Module Map	118
Cross-Reference Table	119
Load Module Format	120
Appendix A. Sample Linkage Editor Programs	122
Sample Program COBFORT	122
Job Control Language	122
Linkage Editor Output	123
Sample Program RPLACJOB	123
Job Control Language	125
Linkage Editor Control Statements	126
Linkage Editor Output	127
Sample Program REGNOVLY	127
Job Control Language	128
Linkage Editor Control Statements	129
Linkage Editor Output	130
Sample Program PARTDS	132
Job Control Language	133
Linkage Editor Control Statements	134
Linkage Editor Output	135
Appendix B. Linkage Editor Requirements and Capacities	136
Capacities	136
Intermediate Data Set	138
Appendix C. Designing and Specifying Overlay Programs	139
Design of an Overlay Program	139
Single Region Overlay Program	140
Control Section Dependency	140
Segment Dependency	142
Length of an Overlay Program	143
Segment Origin	144
Communication between Segments	145
Overlay Process	147
Multiple Region Overlay Program	148
Specification of an Overlay Program	150
Segment Origin	151
Region Origin	152
Positioning Control Sections	153
Using Object Decks	154

Using INCLUDE Statements	154
Using INSERT Statements	155
Special Options	156
OVLY Option	157
LET Option	157
XCAL Option	157
AMODE and RMODE Options	157
Special Considerations	157
Common Areas	157
Storage Requirements	159
Overlay Communication	160
CALL Statement or CALL Macro Instruction	161
Branch Instruction	161
Segment Load (SEGLD) Macro Instruction	162
Segment Wait (SEGWT) Macro Instruction	163
Part II. Loader	165
Chapter 9. Overview and Uses of the Loader	166
Functional Characteristics	166
Compatibility and Restrictions	169
Time Sharing Option (TSO)	169
Processing Object Modules in Virtual Storage	169
Chapter 10. Preparing Input for the Loader	170
Input for the Loader	170
EXEC Statement	170
PARM Field Format	170
Loader Options	171
CALL NOCALL: Automatically Searching SYSLIB	171
EP=name: Specifying the Program Entry Point	171
LET NOLET: Executing with Severity 2 Errors	172
MAP NOMAP: Printing a Program Map	172
NAME=name: Identifying the Loaded Program	172
PRINT NOPRINT: Printing Messages on SYSLOUT	172
RES NORES: Automatically Searching the Link Pack Area Queue	173
SIZE=size: Specifying Virtual Storage	173
TERM NOTERM: Sending Messages to SYSTEMR	173
EXEC Statement Example	173
DD Statements	174
SYSLIN DD Statement	175
SYSLIB DD Statement	175
SYSLOUT DD Statement	176
SYSTEMR DD Statement	176
Loaded Program Data	177
Sample Input for the Loader	177
Chapter 11. Invoking The Loader	180
Chapter 12. Interpreting Loader Output	185
Appendix D. Loader Storage Considerations	187
Appendix E. Loader Return Codes	189
Glossary	191
Index	194

FIGURES

1.	Preparing a Source Module for Execution	4
2.	Preparing a Source Module for Execution, and Executing the Load Module	5
3.	External Names and External References	6
4.	Use of the External Symbol Dictionary	8
5.	Input, Intermediate, and Output Sources for the Linkage Editor	10
6.	A Load Module Produced by the Linkage Editor	11
7.	Linkage Editor Processing—Module Linkage	13
8.	Linkage Editor Processing—Module Editing	14
9.	Linkage Editor Processing—Additional Input Sources	15
10.	System Automatic Call Libraries	29
11.	Processing of One INCLUDE Control Statement	33
12.	Processing of More than One INCLUDE Control Statement	34
13.	SYSUT1 and SYSLMOD Device Types and Their Maximum Record Sizes	46
14.	Load Module Buffer Area and SYSLMOD and SYSUT1 Record Sizes	47
15.	Incompatible Job Step Options for the Linkage Editor	53
16.	Linkage Editor Return Codes	54
17.	Linkage Editor ddnames	56
18.	DCB Requirements for Object Module and Control Statement Input	57
19.	DCB Requirements for SYSPRINT	58
20.	DCB Requirements for Additional Input Data Sets	60
21.	Statements in the LKED Cataloged Procedure	62
22.	Statements in the LKEDG Cataloged Procedure	64
23.	Overlay Structure for INSERT Statement Example	78
24.	Output Load Module for ORDER Statement Example	85
25.	Overlay Structure for OVERLAY Statement Example	87
26.	Output Load Module for PAGE Statement Example	89
27.	Editing a Module	94
28.	Changing an External Reference and an Entry Point	97
29.	Automatic Replacement of Control Sections	100
30.	Replacing a Control Section with the REPLACE Control Statement	102
31.	Deleting a Control Section	103
32.	Ordering Control Sections	104
33.	Aligning Control Sections on Page Boundaries	106
34.	Diagnostic Messages Issued by the Linkage Editor	118
35.	Module Map	119
36.	Cross-Reference Table	120
37.	Load Module Format	121
38.	Linkage Editor Output for Sample Program COBFORT	124
39.	Linkage Editor Output for Job Step that Created SUBONE	125
40.	Job Control Statements for RPLACJOB	126
41.	Linkage Editor Control Statements for RPLACJOB	126
42.	Linkage Editor Output for Sample Program RPLACJOB	127
43.	Overlay Tree for Multiple-Region Sample Program REGNOVLY	128
44.	Job Control Statements for REGNOVLY	129
45.	Linkage Editor Output for Sample Program REGNOVLY	130
46.	Input Statements for IEBUPDTE Utility Program	133
47.	Job Control Statements for PARTDS	134
48.	Linkage Editor Capacities for Minimal SIZE Values (96K bytes, 6K bytes)	136
49.	Control Section Dependencies	141
50.	Single-Region Overlay Tree Structure	142
51.	Length of an Overlay Module	143
52.	Segment Origin and Use of Storage	144
53.	Inclusive and Exclusive Segments	145
54.	Inclusive and Exclusive References	146
55.	Location of Segment and Entry Tables in an Overlay Module	147
56.	Control Sections Used by Several Paths	149
57.	Overlay Tree for Multiple-Region Program	150
58.	Symbolic Segment Origin in Single-Region Program	152

59.	Symbolic Segment and Region Origin in Multiple-Region Program	153
60.	Common Areas before Processing	158
61.	Common Areas after Processing	159
62.	Branch Sequences for Overlay Programs	161
63.	Use of the SEGLD Macro Instruction	162
64.	Use of the SEGWT Macro Instruction	163
65.	Loader Processing—SYSLIB Resolution	167
66.	Loader Processing—Link Pack Area and SYSLIB Resolution	168
67.	Loader Processing—Automatic Editing	168
68.	Input Deck for the Loader—Basic Format	170
69.	Loader and Loaded Program Data Input Stream	177
70.	Input Deck for a Load Job	177
71.	Input Deck for a Compile-Load Job	178
72.	Input Deck for Compilation and Loading of the Three Modules	179
73.	Using the LINK Macro Instruction to Refer to the Loader	181
74.	Using the LOAD and CALL Macro Instructions to Refer to HEWLOADR (Loading Without Identification)	183
75.	Using the LOAD and CALL Macro Instructions to Refer to HEWLOAD (Loading With Identification)	184
76.	Module Map Format Example	186
77.	Virtual Storage Requirements	187
78.	Return Codes	189

INTRODUCTION

The linkage editor and the loader processing programs prepare the output of language translators for execution. The linkage editor prepares a load module that is to be brought into storage for execution by program fetch. The loader prepares the executable program in storage and passes control to it directly.

The linkage editor provides several processing facilities, such as creating overlay programs and aiding program modification. (The linkage editor is also used to build and edit system libraries.) The loader provides high performance loading of programs that do not require the special processing facilities of the linkage editor.

Use of the linkage editor is recommended in the following cases:

- If the program requires linkage editor services in addition to the MAP, LET, NCAL, and SIZE options
- If the program uses linkage editor control statements, such as INCLUDE, NAME, OVERLAY
- If a load module is to be produced for a program library

Use of the loader is recommended if the program only requires the use of the following linkage editor options: MAP, LET, NCAL, and SIZE. Because of its fewer options and because it can process a job in one job step, the loader reduces editing and loading time by about one-half.

Linkage editor processing is performed in a link-edit step. The linkage editor can be used for compile-link edit-go, compile-link edit, link-edit, and link-edit-go jobs. Loader processing is performed in a load step, which is equivalent to the link-edit-go steps. The loader can be used for compile-load and load jobs.

The MVS/370 Data Facility Product linkage editor runs in 24-bit addressing mode.

Details of how each language interfaces with the linkage editor can be found in the publication(s) describing that language.



CHAPTER 1. OVERVIEW

Linkage editor processing is a necessary step that follows the source program assembly or compilation of any problem program. The linkage editor is both a processing program and a service program used in association with the language translators.

Every problem program is designed to fulfill a particular purpose. To achieve that purpose, the program can generally be divided into logical units that perform specific functions. A logical unit of coding that performs a function, or several related functions, is a module. Separate functions should be programmed into separate modules, a process called modular programming. Each module can be written in the symbolic language that best suits the function to be performed. (The symbolic languages are Assembler, ALGOL, BASIC, COBOL, FORTRAN, PASCAL, PL/I, and RPG.)

Each module is separately assembled or compiled by one of the language translators. The input to a language translator is a source module; the output from a language translator is an object module. Before an object module can be executed, it must be processed by the linkage editor. The output of the linkage editor is a load module (Figure 1).




Figure 1. Preparing a Source Module for Execution

An object module is in relocatable format with machine code that is not executable. A load module (see "Load Module Format" on page 120) is also relocatable, but with executable machine code. A load module is in a format that can be loaded into virtual storage and relocated by program fetch (Figure 2 on page 5).

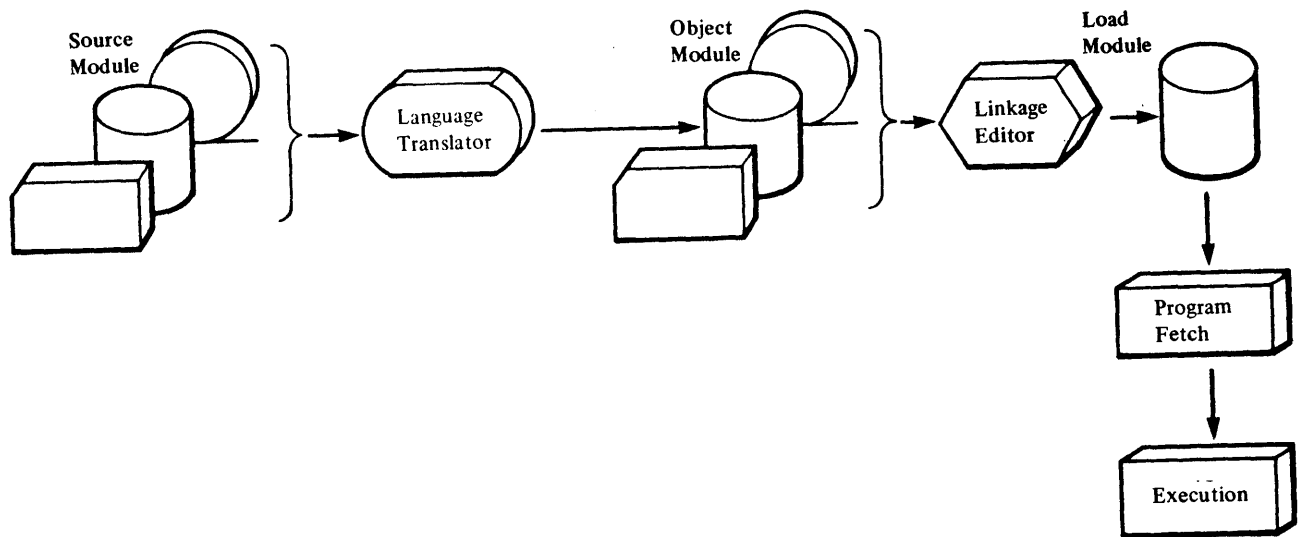


Figure 2. Preparing a Source Module for Execution, and Executing the Load Module

Any module is composed of one or more control sections. A control section is a unit of coding (instructions and data) that is, in itself, an entity. All elements of a control section are loaded and executed in a constant relationship to one another. A control section is, therefore, the smallest separately relocatable unit of a program.

Each module in the input to the linkage editor may contain symbolic references to control sections in other modules; such references are called external references. These references are made by means of address constants (adcons). The symbol referred to by an external reference must be either the name of a control section or the name of an entry point in a control section. Control section names and entry names are called external names. By matching an external reference with an external name, the linkage editor resolves references between modules. External references and external names are called external symbols (Figure 3 on page 6). An external symbol is one that is defined in one module and can be referred to in another.

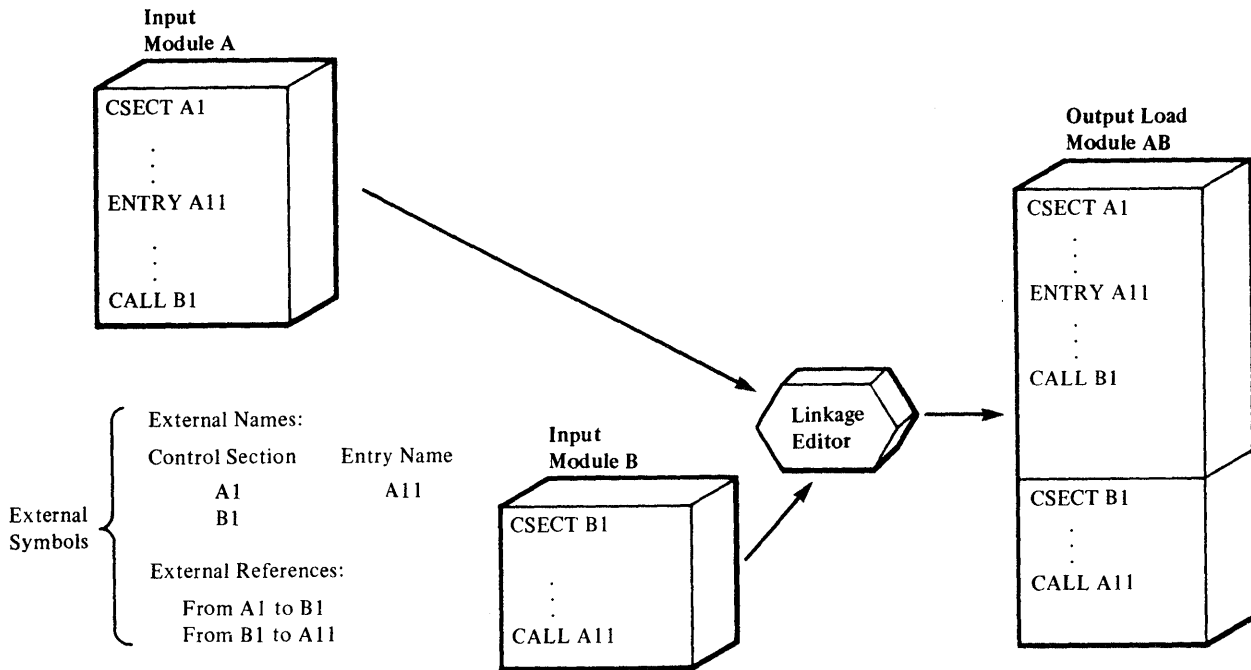


Figure 3. External Names and External References

OBJECT AND LOAD MODULES

Object modules and load modules have the same basic logical structure. Each consists of:

- Control dictionaries, containing the information necessary to resolve symbolic cross-references between control sections of different modules, and to relocate address constants. Control dictionary entries are generated when external symbols, address constants, or control sections are processed by a language translator. Each language translator usually produces two kinds of control dictionaries: an external symbol dictionary (ESD) and a relocation dictionary (RLD).
- Text, containing the instructions and data of the program.
- An end-of-module indication: an END statement in an object module, an end-of-module indicator in a load module.

Each control dictionary, text, and end indication is described in greater detail below.

Both object modules and load modules can contain data used by the linkage editor to create CSECT identification records (IDR). If the language translator creating an object module supports CSECT identification, the input object module can contain translator data for identification records on the END statement. Input load modules differ from object modules in the type of data they supply. Input load modules can also provide HMASPZAP data, linkage editor data, and user data to the identification records that are built during linkage editor processing. During the link-edit step, the optional IDENTIFY control statement is used to supply the optional user data for the CSECT

identification records. See "IDENTIFY Statement" on page 74 for more information.

The design intent of the Linkage Editor is that object and load modules that can be correctly processed by a previous MVS Linkage Editor will be correctly processed by the MVS/XA Version 2 Linkage Editor.

External Symbol Dictionary

The external symbol dictionary (ESD) contains one entry for each external symbol defined or referred to within a module. The dictionary contains an entry for each external reference, pseudoregister (external dummy section), entry name, named or unnamed control section, and blank or named common area. An entry name, pseudoregister, or named control section can be referred to by any control section or separately processed module; an unnamed control section cannot.

Each entry identifies a symbol, or a symbol reference, and gives its location, if known, within the module. Each entry in the external symbol dictionary is classified as one of the following:

- External reference—a symbol that is defined as an external name in another separately processed module, but is referred to in the module being processed. The external symbol dictionary entry specifies the symbol; the location is unknown.
- Weak external reference—a special type of external reference that is not to be resolved by automatic library call unless an ordinary external reference to the same symbol is found. The external symbol dictionary entry specifies the symbol; the location is unknown.
- Entry name—a name that defines an entry point within a control section. The external symbol dictionary entry specifies the symbol and its location, and identifies the control section to which it belongs.
- Control section name—the symbolic name of a control section. The external symbol dictionary entry specifies the symbol, the length of the control section, and its location. In this case, the location represents the origin of the control section, which is the first byte of the control section. This external symbol dictionary entry specifies the addressing mode and residence mode of the control section, and whether the control section is read-only.
- Blank or named common area—a control section used to reserve a virtual storage area that can be referred to by other modules. The reserved storage area can be used, for example, as a communications region within a program or to hold data supplied at execution time. The external symbol dictionary entry specifies the name, if there is one, and the length of the area. If there is no name, the name field contains blanks.
- Private code—an unnamed control section. This external symbol dictionary entry specifies the length of the control section and the origin. The name field contains blanks. The external symbol dictionary entry may also specify the addressing mode and residence mode of the control section and whether or not the control section is read-only.
- Pseudoregister—a special facility (corresponding to the external dummy section feature of Assembler H Version 2) that can be used to write reenterable programs. A pseudoregister is a dynamically obtained word in virtual storage that can be used as a pointer to dynamically acquired storage; that is, the space for such areas is not reserved in the load module but is acquired during

execution. The external symbol dictionary contains the name, length, alignment, and displacement of the pseudoregister.

When processing input modules, the linkage editor resolves references between modules by matching the referenced symbols to defined symbols. To do this, the linkage editor searches for the external symbol definition in the external symbol dictionary of each input module. As shown in Figure 4, the linkage editor matches the external reference to B1 by locating the definition for B1 in the external symbol dictionary of Module B. In the same way, it matches the external reference to A11 by locating the definition for A11 in the external symbol dictionary of Module A.

Note: External names, including CSECT names and entry names, must be 1 to 8 alphanumeric characters in length. No leading or embedded blanks are permitted, nor are the following characters permitted:

, (or)

All other characters in the 48-character set are permitted in any character position of the name by the linkage editor, including:

+ - = . * ' / &

Special characters should be used with caution, however, because the compilers and assemblers that produce the object decks usually have a more limited character set.

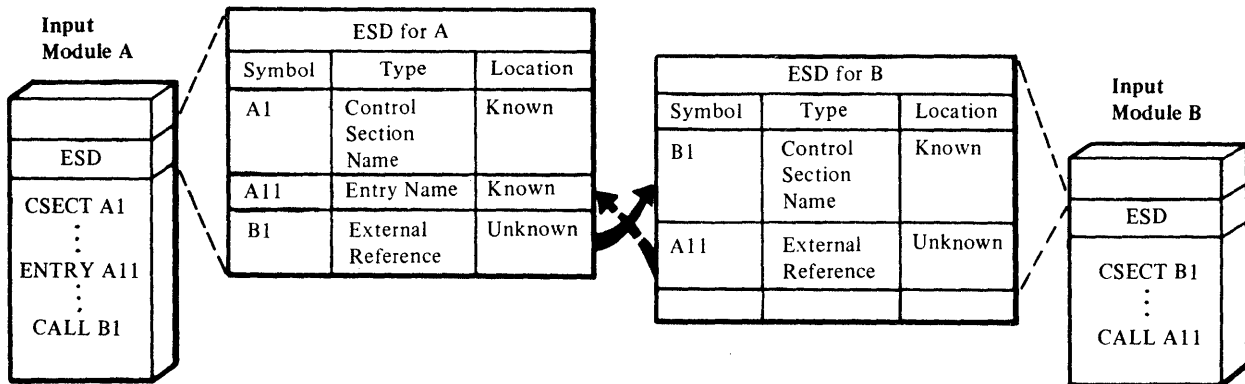


Figure 4. Use of the External Symbol Dictionary

Text

The text contains the instructions and data of the module.

Note: Object module text records may not necessarily be in ascending address sequence (it is possible that the language translator may have created them out of order). When processing large object modules with out-of-order text, the performance of the linkage editor may be improved by presorting the object module text in ascending address sequence (columns 6 through 8 of the text record).

Relocation Dictionary

The relocation dictionary (RLD) contains one entry for each relocatable address constant that must be modified before a module is executed. An entry identifies an address constant by indicating both its location within a control section and the external symbol whose value must be used to compute the value of the address constant. (The external symbol is defined in an external symbol dictionary entry in another control section or module.)

The linkage editor uses the relocation dictionary whenever it processes a module to adjust the address constants for references to other control sections and modules. This dictionary is also used to adjust these address constants again after program fetch reads an output load module from a library and loads it into virtual storage for execution.

End Indication

The end of a load module is marked by an end-of-module indicator (EOM). The EOM cannot, unlike the assembler END instruction, specify an entry point. Therefore, whenever a load module is reprocessed by the linkage editor, a main entry point should be specified on an ENTRY statement. If one is not specified, the linkage editor will assign the first byte of the first control section encountered as the entry point. The programmer will not usually be concerned with the format of records in the object deck. Record formats are described in the appendix of Linkage Editor Logic.

LINKAGE EDITOR PROCESSING

This section discusses the input and output sources of the linkage editor, and the way in which the linkage editor produces a load module.

INPUT AND OUTPUT SOURCES

The linkage editor accepts two major types of input:

- Primary input, which can contain only object modules and linkage editor control statements (called control statements in the following text).
- Additional user-specified input, which can contain either object modules and control statements, or load modules. This input is either specified by the user as input, or incorporated automatically by the linkage editor from a call library.

During processing, the linkage editor generates intermediate data. Intermediate data is placed on a direct access storage device when virtual storage allocated for input data is exhausted.

Output of the linkage editor is of two types:

- A load module, which is always placed in a library (a partitioned data set) as a named member
- Diagnostic output, which is produced as a sequential data set

Figure 5 on page 10 shows the input, intermediate, and output sources for the linkage editor program.

LOAD MODULE CREATION

In processing object and load modules, the linkage editor assigns consecutive relative virtual storage addresses to all control sections and resolves all references between control sections. Object modules produced by several different language translators can be used to form one load module.

An output load module is composed of all input object modules and input load modules processed by the linkage editor. The control dictionaries of an output module are, therefore, a composite of all the control dictionaries in the linkage editor input. The control dictionaries of a load module are called the composite external symbol dictionary (CESD) and the relocation dictionary (RLD). The load module also contains all of the text from each input module, and one end-of-module indicator (see Figure 6 on page 11). See also "Load Module Format" on page 120 for the format of a load module.




Figure 5. Input, Intermediate, and Output Sources for the Linkage Editor

Assigning Addresses

Each module to be processed by the linkage editor has an origin that was assigned during assembly, compilation, or a previous execution of the linkage editor. When several modules, each with an independently assigned origin, are to be processed by the linkage editor, the sequence of the addresses is unpredictable; two input modules may even have the same origin.

Each input module can be made up of one or more control sections. To produce an executable output load module, the linkage editor assigns relative virtual storage addresses to each control section by assigning an origin to the first control section encountered and then assigning addresses, relative to that origin, to all other control sections to be included in the output load module. The value assigned as the origin of the control section is used to relocate each address-dependent item in the control section.

Although the addresses in a load module are consecutive, they are all relative to base zero. When a load module is to be executed, program fetch prepares the module for execution by loading it at a specific virtual storage location. The addresses in the module are then increased by this base address. Each address constant must also be readjusted, another function of program fetch.




Figure 6. A Load Module Produced by the Linkage Editor

Resolving External References

The linkage editor also resolves external references in input modules. Cross-references between control sections in different modules are symbolic. They must be resolved relative to the addresses assigned to the load module. The linkage editor calculates the new address of each relocatable expression in a control section and determines the assigned origin of the item to which it refers.

CHAPTER 2. USES OF THE LINKAGE EDITOR

LINKAGE EDITOR INPUT

Linkage editor input may consist of a combination of object modules, load modules, and control statements. The primary function of the linkage editor is to combine these modules, in accordance with the requirements stated on control statements, into a single output load module. Although this linking or combining of modules is its primary function, the linkage editor also:

- Edits modules by replacing, deleting, rearranging, and ordering control sections as directed by control statements
- Aligns control sections and named common areas on 4K-byte page boundaries as directed by control statements
- Accepts additional input modules from data sets other than the primary input data set, either automatically or upon request
- Reserves storage for the common control sections generated by Assembler and FORTRAN language translators, and static external areas generated by PL/I
- Computes total length and assigns displacements for all pseudoregisters (external dummy sections)
- Creates overlay programs in a structure defined by control statements
- Creates multiple output load modules as directed by control statements
- Provides special processing and diagnostic output options
- Assigns module attributes that describe the structure, content, and logical format of the output load module
- Allocates storage areas for linkage editor processing as specified by the programmer
- Stores system status index information in the directory of the output module library (systems personnel only)
- Traces the processing history of a program
- Allows the user to lengthen a control section or named common section without changing source code, reassembling, or recompiling
- Allows the user to assign an authorization code to a load module that (a) makes it a restricted resource and (b) enables it to pass control to other restricted resources
- Assigns an addressing mode for the main entry point, all true aliases, and each alternate entry point into the output load module
- Assigns a residence mode for the output load module
- Indicates which control sections are read-only (relevant only in creating a nucleus load module for MVS/XA)

Each of the linkage editor functions is described in the following paragraphs.

Links Modules

Processing by the linkage editor makes it possible for the programmer to divide a program into several modules, which can be separately assembled or compiled, and each containing one or more control sections. The linkage editor combines these modules into one output load module (see Figure 7) with contiguous, virtual storage addresses. During processing by the linkage editor, references between modules within the input are resolved. The output module is placed in a library (partitioned data set).

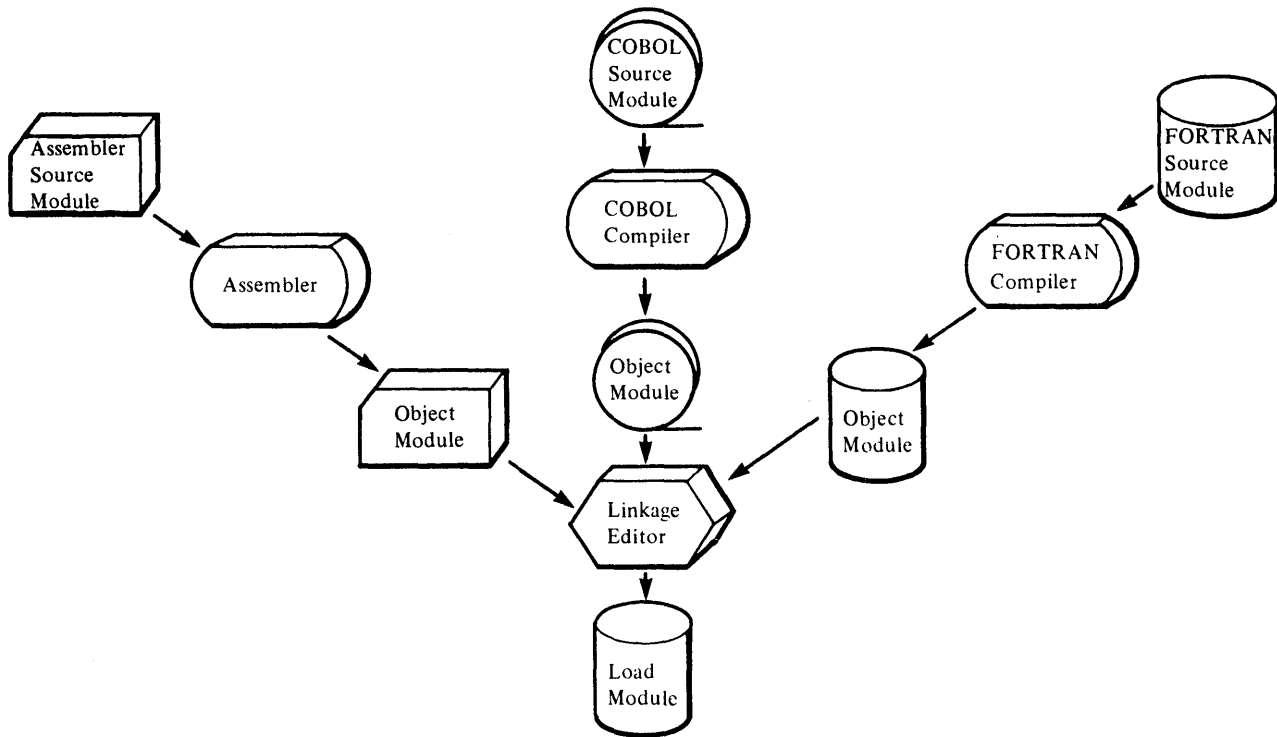


Figure 7. Linkage Editor Processing—Module Linkage

Edits Modules

Program modification is made easier by the editing functions of the linkage editor. When the functions of a program are changed, the programmer modifies, then compiles and link-edits again, only the affected control sections instead of the entire source module.

Control sections can be replaced, renamed, deleted, moved, or ordered as directed by control statements. Control sections can also be automatically replaced by the linkage editor. External symbols can be changed or deleted as directed by control statements.

Figure 8 on page 14 illustrates the module editing function of the linkage editor.

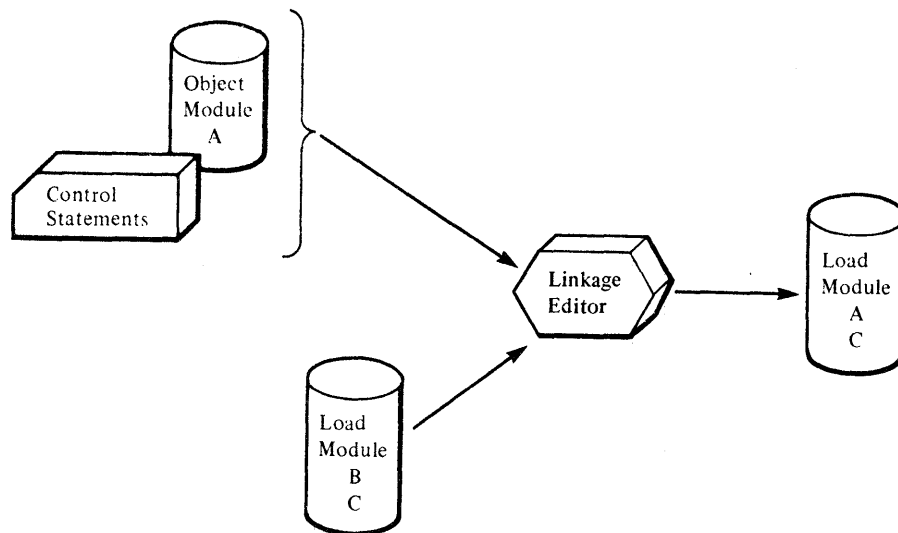


Figure 8. Linkage Editor Processing—Module Editing

Aligns Control Sections or Common Areas on Page Boundaries

Control sections or named common areas in the output load module can be aligned on 4K-byte page boundaries. Alignment on page boundaries enables the programmer to use real storage more efficiently and thus appreciably reduce the paging rate for the job.

Accepts Additional Input Sources

Standard subroutines can be included in the output module, thus reducing the work in coding programs. The programmer can specify that a subroutine be included at a particular time during the processing of the program by using a control statement. When the linkage editor processes a program that contains this statement, the module containing the subroutine is retrieved from the indicated input source and made a part of the output module (Figure 9 on page 15).

Symbols that are still undefined after all input modules have been processed cause the automatic library-call mechanism to search for modules that will resolve these references. When a module name is found that matches the unresolved symbol, the module is processed by the linkage editor and also becomes part of the output module (Figure 9).

Note: The linkage editor distinguishes a special type of external reference—the weak external reference. An unresolved weak external reference does not cause the linkage editor to use the automatic library-call mechanism. Instead, the reference is left unresolved, and the load module is marked as executable.

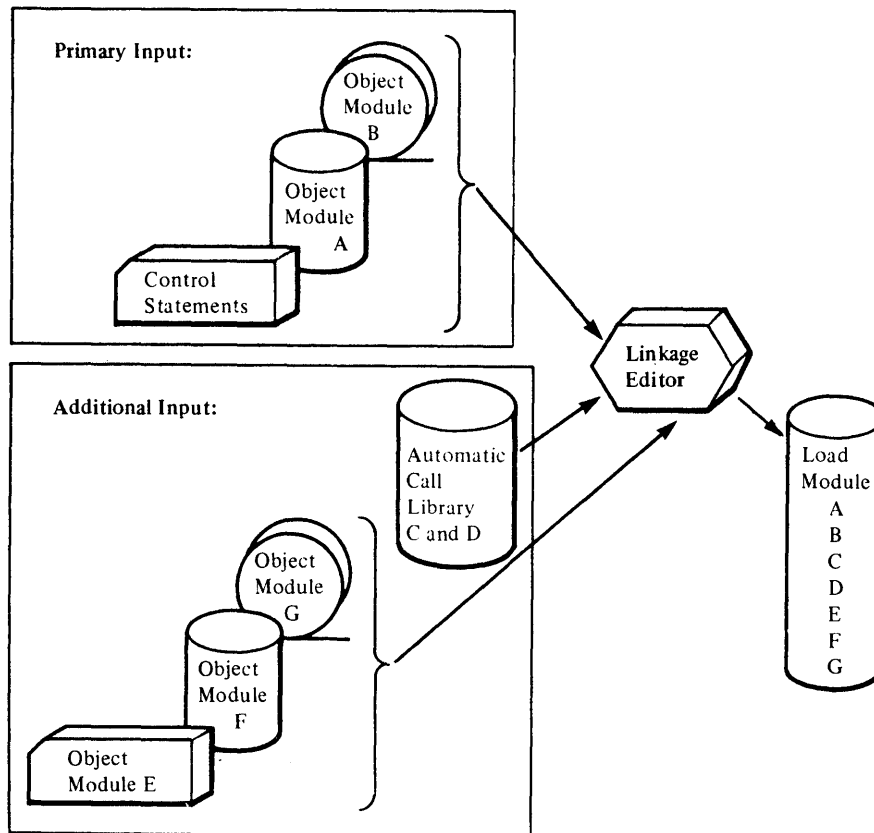


Figure 9. Linkage Editor Processing—Additional Input Sources

Reserves Storage

The linkage editor processes common control sections generated by the FORTRAN and Assembler language translators. The static external storage areas generated by the PL/I compiler are processed in the same way. The common areas are collected by the linkage editor, and a reserved virtual storage area is provided within the output module.

Processes Pseudoregisters

Pseudoregisters, like the external dummy sections of Assembler H Version 2, aid in generating reenterable code. The linkage editor processes pseudoregisters by accumulating the total length of storage required for all pseudoregisters and recording the displacement of each. During execution, the program dynamically acquires the necessary storage.

Creates Overlay Programs

To minimize virtual storage requirements, the programmer can organize a program into an overlay structure by dividing it into segments according to the functional relationships of the control sections. Two or more segments that need not be in virtual storage at the same time can be assigned the same relative virtual storage addresses, and can be loaded at different times.

The programmer uses control statements to specify the relationship of segments within the overlay structure. The segments of the load module are placed in a library so that the control program can load them separately when the load module is executed.

Creates Multiple Load Modules

The linkage editor can also process its input to form more than one load module within a single job step. Each load module is placed in the library under a unique member name, as specified by a control statement.

Provides Special Processing and Diagnostic Output Options

The programmer can specify special processing options that negate automatic library call or the effect of minor errors. In addition, the linkage editor can produce a module map or cross-reference table that shows the arrangement of control sections in the output module and indicates how they communicate with one another. A list of the control statements processed can also be produced.

Throughout processing, errors and possible error conditions are logged. Serious errors cause the linkage editor to mark the output module not executable. Additional diagnostic data is automatically logged by the linkage editor. The data indicates the disposition of the load module in the output module library.

Assigns Load Module Attributes

When the linkage editor generates a load module, it places an entry for the module in the directory of the library. This entry contains attributes that describe the structure, content, and logical format of the load module. The control program uses these attributes to determine how a module is to be loaded, what it contains, if it is executable, whether it is executable more than once without reloading, and if it can be executed by concurrent tasks. Some module attributes can be specified by the programmer; others are specified by the linkage editor as a result of information gathered during processing. See also "Assigns Addressing Mode" on page 18, "Assigns Residence Mode" on page 19, and "Assigns Read-only Attribute" on page 21.

Allocates User-Specified Virtual Storage Areas

The programmer can specify the total amount of virtual storage to be made available to the linkage editor, the amount to be used for the load module buffer, and the buffer for the output load module.

Stores System Status Index Information

The following information is intended for systems personnel responsible for maintaining IBM-supplied load modules. It is not generally applicable to non-IBM load modules.

Four bytes in the library directory entry for IBM-supplied load modules are used to store system status index information. This information, which is used for maintenance of the modules, is placed in the directory with a control statement.

Traces Processing History

Tracing the processing history of a program is simplified by the CSECT identification (IDR) records created and maintained by the linkage editor. A CSECT identification record can contain data that describes:

- The language translator, its level, and the translation date for each control section
- The most recent processing by the linkage editor
- Any modification made to the executable code of any control section

Optionally, user-supplied data associated with the executable code of a control section can also be recorded.

Lengthens Control Sections or Named Common Sections

The user can lengthen control sections or named common sections of a program to add patch space without changing the source code, reassembling, or recompiling.

Added space, consisting of binary zeros, is put at the end of a specified control section by using the EXPAND control statement (see "Chapter 5. Specifying an Operation with Control Statements" on page 67). Space cannot be added to a private code or blank common section.

Assigns an Authorization Code to Output Load Modules

The authorized program facility (APF) limits the use of sensitive system and (optionally) user services and resources to authorized system and user programs. Authorization is defined as access to those services and resources. The services and resources to which access is limited are described in Initialization and Tuning Guide.

Programs are authorized at the job-step level. For a job step to gain authorization initially, the first module loaded at the start of the job step must be an authorized module, and it must have been loaded from an authorized library. Otherwise, the job step is not authorized initially and cannot subsequently gain authorization.

For a job step to maintain its authorization, all subsequent modules invoked during the job step (via LINK, LOAD, ATTACH, and/or XCTL macro instructions) must be loaded from an authorized library. Otherwise, the job step loses its authorization and cannot regain authorization.

A library becomes an "authorized" library by the inclusion of its name in a list called IEAAPF00. This list is described in more detail in Initialization and Tuning Guide.

A load module becomes "authorized" by the assignment of an authorization code to the load module during linkage-editing. This assignment is made via the PARM field parameter AC or via

the control statement SETCODE, which are described in the sections that follow. See "SETCODE Statement" on page 92.

Assigns Addressing Mode

The addressing mode (AMODE) is the attribute of an entry point into a load module that specifies the addressing mode in effect when the load module is entered at that entry point at execution time.

The valid addressing modes are:

- 24 Indicating that 24-bit addressing will be in effect
- 31 Indicating that 31-bit addressing will be in effect
- ANY Indicating that either 24-bit or 31-bit addressing may be in effect

The linkage editor determines the addressing mode for an entry point (either the main entry point, its true alias, or an alternate entry point) according to the following rules:

- The linkage editor assigns a default AMODE of 24. This is done only in the absence of a valid, explicit specification of the addressing mode for the entry point.
- The linkage editor assigns the AMODE values contained in the object module's ESD. These AMODE values were specified by the user at assembly time and represent the AMODE values assigned to the entry points within the CSECTs and private code for the module.
- The linkage editor assigns all the entry points into the load module (the main entry point, its true aliases, and the alternate entry points) the AMODE value specified as a parameter in the PARM field of the EXEC statement. This AMODE value overrides the AMODE value, if any, found in the ESD data.
- The linkage editor assigns the AMODE value specified as an operand on the MODE control statement to all of the entry points into the load module (the main entry point, its true aliases, and the alternate entry points). This AMODE value overrides any value specified as a parameter in the EXEC statement or any values found in the ESD data.

The linkage editor provides the AMODE value for each entry point into the load module in its directory entry. In the case of a true alias of the main entry point or an alternate entry point, the directory entry contains the AMODE value for both the alias/alternate entry point and the main entry point.

The AMODE value provided to the linkage editor in the ESD data of an object module is retained in the ESD data of the load module, for use in subsequent link-editing, except in the case of a load module built for overlay. In building a load module for overlay, the AMODE value in the ESD data of the load module is lost and can only be reintroduced by inclusion of the object module(s) carrying that value. Use of the overriding AMODE specifications (the parameter in the PARM field of the EXEC statement or the operand in the MODE control statement) establishes the AMODE value carried in the directory entry, but does not affect the ESD data.

All entry points in load modules built for overlay are assigned an AMODE of 24, regardless of the ESD data, the PARM field parameter, or the MODE statement operand.

Assigns Residence Mode

The residence mode (RMODE) is the attribute of a load module that specifies the residence mode of a load module when it is loaded into virtual storage for execution.

The valid residence modes are:

- 24 Indicating that the module must reside within 24-bit addressable virtual storage (that is, below the 16-megabyte virtual storage line)
- ANY Indicating that the module may reside anywhere in virtual storage (that is, either above or below the 16-megabyte virtual storage line)

The linkage editor determines the residence mode for a load module according to the following rules:

- The linkage editor assigns a default RMODE of 24. This occurs only in the absence of a valid explicit specification of the residence mode for the load module.
- The linkage editor assigns the RMODE specified in the object module. This RMODE value is specified by the user to the assembler for the control section or private code. The RMODE value passes to the linkage editor in the ESD data. The linkage editor assigns the RMODE value taken from the control section or private code that contributes to the output load module, ignoring identically named control sections and private code that are replaced or deleted.
- As the control sections and private code that contribute to the output load module are processed, the RMODE value for the load module, based on the ESD data, is accumulated on a "most restrictive" basis. This means:
 - If any section in the load module has an RMODE of 24, the RMODE for the load module is 24.
 - If all sections in the load module have an RMODE of ANY, the RMODE for the load module is ANY.
- The linkage editor assigns to the load module the RMODE value specified as a parameter in the PARM field of the EXEC statement. This RMODE value overrides the RMODE value, if any, found in the ESD data.
- The linkage editor assigns to the load module the RMODE value specified as an operand on the MODE control statement. This RMODE value overrides the RMODE value, if any, specified as a parameter in the PARM field of the EXEC statement as well as the RMODE value, if any, found in the ESD data.

Load modules built for overlay are assigned an RMODE of 24, regardless of the ESD data, the PARM field parameter, or the MODE statement operand.

The linkage editor provides the RMODE value for the load module in each directory entry applicable to that load module.

Except in the case of a load module built for overlay, the RMODE value provided to the linkage editor in the ESD data of an object module is retained in the ESD data of the load module, for use in subsequent link-editing. In building a load module for overlay, the RMODE value in the ESD data of the load module is lost and can only be reintroduced by inclusion of the object module(s) carrying that value. Use of the overriding RMODE specifications (the parameter in the PARM field of the EXEC statement or the operand in the MODE control statement) establishes the RMODE value carried in the directory entry, but does not affect the ESD data.

AMODE/RMODE Combinations from the ESD

When AMODE and RMODE data have not been specified on either a MODE linkage editor control statement or in the PARM field of the EXEC statement, the linkage editor determines the AMODE for each entry point and the RMODE for the load module based on ESD data. (Load module entry point designation is discussed under "Entry Point" on page 112.) The linkage editor validates the six possible AMODE/RMODE combinations from the ESD as follows:

	RMODE=24	RMODE=ANY
AMODE=24	valid	invalid
AMODE=31	valid	valid
AMODE=ANY	valid	invalid

Load module entry points (main and alternate) may be either control section name external symbols or entry name external symbols.¹ (See "External Symbol Dictionary," the section on Control section name on page 7.) When an entry point is a control section name, the linkage editor acquires AMODE and RMODE data directly from the control section name ESD entry. When an entry point is an entry name external symbol, the linkage editor acquires AMODE and RMODE data from the associated control section name ESD entry.

Based on the AMODE and RMODE data acquired from the ESD, the linkage editor determines a load module RMODE (see "Assigns Residence Mode" on page 19), and assigns an AMODE to each entry point as outlined below:

- If an entry point external symbol is marked with any of the allowable AMODE values and an RMODE of 24, the entry point is assigned the same AMODE attribute as its associated external symbol.
- The AMODE 24/RMODE ANY combination is invalid as it could allow 24-bit addressing above the 16Mb line. The linkage editor should never find this combination in the ESD since it is flagged by IBM compilers and assemblers as an error condition. If it does find this combination, the linkage editor issues a non-terminal error message, forces the load module RMODE to 24, and assigns an AMODE of 24 to the entry point.
- If the entry point external symbol is marked AMODE 31/RMODE ANY, the entry point AMODE will be 31 and the RMODE will be that of the load module.
- If the entry point external symbol is marked AMODE ANY/RMODE ANY, associated entry point attributes are assigned according to the following hierarchy:
 - If the load module contains one or more CSECTs marked AMODE 24, the linkage editor assigns an AMODE of 24 to all entry points that have ESD entries marked AMODE ANY/RMODE ANY.
 - If the load module has an RMODE of 24 and it contains no CSECTs marked AMODE 24, the linkage editor assigns an AMODE of ANY to these entry points.

¹ The main entry point to a load module is usually an external symbol, although when specified on an assembler language END statement, it may be a displacement into the CSECT. Alternate entry points must always be external symbols.

- If the load module RMODE is ANY, the linkage editor assigns an AMODE of 31 to these entry points.

AMODE/RMODE Hierarchy

The following hierarchy is used to determine the addressing and residence modes of the linkage editor output:

1. Value on the linkage editor MODE statement
2. Value of the parm field on the EXECUTE statement
3. Value in the ESD data produced by the AMODE= or RMODE= assembler statement
4. Default value of 24

Note: An overlay module always results in an AMODE of 24 and an RMODE of 24. A load module produced from multiple object modules results in an RMODE of 24, if any one of the object modules has an RMODE of 24.

Assigns Read-only Attribute

A read-only control section (RSECT) is defined by the user in the source language which assembles the control section. The assembler indicates in the external symbol dictionary entry for the control section that it is read-only. The linkage editor reflects that indication in the scatter table for the output load module.

The indication of the read-only attribute is relevant only to the nucleus initialization program in MVS/XA. In all other cases it is ignored.

RELATIONSHIP TO THE OPERATING SYSTEM

The linkage editor has the same relationship to the operating system as any other processing program. It can be executed either as a job step, a subprogram, or a subtask. Control is passed to the linkage editor in one of three ways:

- As a job step, when the linkage editor is specified on an EXEC job control statement in the input stream
- As a subprogram, with the execution of a CALL macro instruction (after the execution of a LOAD macro instruction), a LINK macro instruction, or an XCTL macro instruction
- As a subtask, in multitasking systems, with the execution of the ATTACH macro instruction

Execution of the linkage editor and the data sets used by the linkage editor are described to the system with job control language statements. These statements describe all jobs to be performed by the system.

Note: Job control statements should not be confused with linkage editor control statements. Job control statements are processed before the linkage editor is executed; linkage editor control statements are processed during linkage editor execution.

Time Sharing Option (TSO)

When the linkage editor is used under TSO, it is invoked by the linkage editor prompter program that acts as an interface between the user, the operating system, and the linkage editor. Under TSO, execution of the linkage editor and definition of data sets used by the linkage editor are described to the system through use of the LINK command that causes the prompter to be executed. Operands of the LINK command can also be used to specify the linkage editor options a job requires. Complete procedures for use of the LINK command are given in TSO Command Language Reference.

CHAPTER 3. DEFINING INPUT TO THE LINKAGE EDITOR

The linkage editor accepts input from two major sources: the primary input data set and additional data sets. The primary input data set is made available through job control statements. Additional data sets are made available either through the automatic library call mechanism, or through user-specified control statements. They must, however, also be defined with job control statements.

Primary and additional input data sets may contain the following types of data:

- One or more object modules
- One or more load modules
- Control statements
- Combinations of the above (restrictions on certain combinations are noted where they apply)

Object modules and control statements may be contained in either sequential or partitioned data sets. Load modules must be contained in partitioned data sets.

This chapter describes the "linking" functions of the linkage editor only; the "editing" functions are described in "Chapter 6. Editing a Control Section" on page 94.

PRIMARY INPUT DATA SET

The primary input data set is required for every linkage editor job step. It must be defined by a DD statement with the ddname SYSLIN. The primary input can be:

- A sequential data set
- A member of a partitioned data set
- A concatenation of sequential data sets and/or members of partitioned data sets

The primary input data set must contain object modules and/or control statements. The modules and control statements are processed sequentially and their order determines the basic order of linkage editor processing during a given execution. However, the order of the control sections after processing does not necessarily reflect the order in which they appeared in the input.

In the examples that follow, only the statements necessary to define the input to the linkage editor are shown; complete examples are shown in "Appendix A. Sample Linkage Editor Programs" on page 122.

OBJECT MODULES

The primary input to the linkage editor may consist solely of one or more object modules. The rest of this section discusses object module input from cards, as a member of a partitioned data set, passed from a previous job step, or created in a separate job.

From Cards

Object module input to the linkage editor may be on cards. The card deck itself is treated as a sequential data set; the cards are placed in the input stream, after a DD * statement, as follows:

```
//SYSLIN DD *
Object Deck A
Object Deck B
/*
```

The card input is followed by a /* statement.

An example of the JCL when card decks are used in addition to other input is as follows:

```
//SYSLIN DD DSNAME=INPUT,...
// DD *
Object Deck A
Object Deck B
/*
```

By omitting the dname on the second DD statement, the card input is concatenated to the data set described on the SYSLIN DD statement.

As a Member of a Partitioned Data Set

An object module in a partitioned data set can be used as primary input to the linkage editor by specifying its data set name and member name on the SYSLIN DD statement. In the following example, the member named TAXCOMP in the object module library LIBROUT is to be the primary input; LIBROUT is a cataloged data set:

```
//SYSLIN DD DSNAME=LIBROUT(TAXCOMP),
// DISP=(OLD,KEEP)
```

The library member is processed as if it were a sequential data set.

Members of partitioned data sets can be concatenated with other input data sets, as follows:

```
//SYSLIN DD DSNAME=OBJLIB,DISP=(OLD,KEEP),...
// DD DSNAME=LIBROUT(TAXCOMP),
// DISP=(OLD,KEEP)
```

Library member TAXCOMP is concatenated to data set OBJLIB; because they are the primary input, both must contain object modules.

Done - things look OK.

nel

Passed from a Previous Job Step

An object module to be used as input can be passed from a previous job step to a linkage editor job step in the same job, as in a compile-link-edit job. That is, the output from the compiler is direct input to the linkage editor. In the following example, an object module that was created in a previous job step (STEPSA) is passed to the linkage editor job step (STEPB):

```
→ STEPSA
//SYSGO DD DSNAME=&&OBJECT,DISP=(NEW,PASS),...
.
.
→ STEPB
//SYSLIN DD DSNAME=&&OBJECT,DISP=(OLD,DELETE)
           (no SPACE parameter)
```

The data set name &&OBJECT, used in both job steps, identifies the object module as the output of the language processor on the SYSGO DD statement, and as the primary input to the linkage editor on the SYSLIN DD statement.

Note: The double ampersand (&&) in the data set name defines a temporary data set. These data sets exist for the duration of the job and are automatically deleted at the end of the job. If the data set is to be preserved for longer than the duration of a single job, the double ampersand is not used (DSNAME=OBJECT).

The method used in the preceding example can also be used to retrieve object modules created in previous steps. If the same data set name is used for the output of each language processor, one SYSLIN DD statement can be used to retrieve all the object modules, as follows:

```
STEPSA:
//SYSGO DD DSNAME=&&OBJMOD,DISP=(NEW,PASS),...
.
.

STEPB:
//SYSPUNCH DD DSNAME=&&OBJMOD,DISP=(MOD,PASS)
.
.

STEPC:
//SYSLIN DD DSNAME=&&OBJMOD,DISP=(OLD,DELETE)
```

The two object modules from STEPSA and STEPB are placed in the same sequential data set, &&OBJMOD. The SYSLIN DD statement in STEPC causes both object modules to be used as the primary input to the linkage editor.

Another method can be used to accomplish this purpose: concatenation of data sets. This method could be used if the object modules were created in previous job steps with different member names, as follows:


```

STEP A:
//SYSGO      DD      DSNAME=&&OBJLIB(MODA),DISP=(NEW,
//            PASS),...
            .
            .
STEP B:
//SYSPUNCH   DD      DSNAME=&&OBJLIB(MODB),DISP=(MOD,
//            PASS),...
            .
            .
STEP C:
//SYSLIN     DD      DSNAME=&&OBJLIB(MODA),DISP=(OLD,
//            DELETE)
//           DD      DSNAME=&&OBJLIB(MODB),DISP=(OLD,
//            DELETE),VOL=REF=*.STEPB.SYSPUNCH

```

The object modules created in STEPA and STEP B were placed in a partitioned data set with different member names. The two members are concatenated in STEPC as primary input. Each member is considered to be a sequential data set.

Created in a Separate Job

If the only input to the linkage editor is an object module from a previous job, the SYSLIN DD statement contains all the information necessary to locate the object module, as follows:

```

//SYSLIN     DD      DSNAME=OBJECT,DISP=(OLD,DELETE),
//            UNIT=3380,VOLUME=SER=LIB613

```

An object module created in a separate job may also be on cards, in which case it is handled as described earlier.

CONTROL STATEMENTS

The primary input data set may also consist solely of control statements. When the primary input is control statements, input modules are specified on INCLUDE control statements (see "Included Data Sets" on page 33). The control statements may be either placed in the input stream or stored in a permanent data set.

In the following example, the primary input consists of control statements in the input stream:

```

//SYSLIN     DD      *
Linkage Editor Control Statements
/*

```

In the next example, the primary input consists of control statements stored in the member INCLUDES in the partitioned data set CTLSTMTS:

```
//SYSLIN    DD    DSNAME=CTLSTMTS(INCLUDES),DISP=(OLD,
//          KEEP),...
```

In either case, the control statements can be any of those described in "Chapter 5. Specifying an Operation with Control Statements" on page 67, as long as the rules given there are followed.

OBJECT MODULES AND CONTROL STATEMENTS

The primary input to the linkage editor may contain both object modules and control statements. The object modules and control statements may be in either the same data set or in different data sets. If the modules and statements are in the same data set, this data set is described on the SYSLIN DD statement as any data set is described.

If the modules and statements are in different data sets, the data sets are concatenated. The control statements may be defined either in the input stream or as a separate data set.

Control Statements in the Input Stream

Control statements can be placed in the input stream and concatenated to an object module data set, as follows:

```
//SYSLIN    DD    DSNAME=&&OBJECT,...
//          DD    *
Linkage Editor Control Statements
/*
```

Another method of handling control statements in the input stream is to use the DDNAME parameter, as follows:

```
//SYSLIN    DD    DSNAME=&&OBJECT,...
//          DD    DDNAME=SYSIN
          .
          .
//SYSIN     DD    *
Linkage Editor Control Statements
/*
```

Note: The linkage editor cataloged procedures use DDNAME=SYSIN for the SYSLIN DD statement to allow the programmer to specify the primary input data set required.

Control Statements in a Separate Data Set

A separate data set that contains control statements may be concatenated to a data set that contains an object module. The control statements for a frequently used procedure (for example, a complex overlay structure or a series of INCLUDE statements) can be stored permanently. In the following example, the members of data set CTLSTMTS contain linkage editor control statements. One of the members is concatenated to data set &&OBJECT.

```
//SYSLIN      DD      DSNAME=&&OBJECT,DISP=(OLD,DELETE),...  
//           DD      DSNAME=CTLSTMTS(OVLY),DISP=(OLD,  
//           KEEP),...
```

The control statements in the member named OVLY of the partitioned data set CTLSTMTS are used to structure the object module.

AUTOMATIC LIBRARY CALL

The automatic library-call mechanism is used to resolve external references that were not resolved during primary input processing. Unresolved external references found in modules from additional data sources are also processed by this mechanism.

Note: The following discussion of automatic library call does not apply to unresolved weak external references; they are left unresolved.

The automatic library-call mechanism involves a search of the directory of the automatic call library for an entry that matches the unresolved external reference. When a match is found, the entire member is processed as input to the linkage editor.

Automatic library call can resolve an external reference when the following conditions exist: The external reference must be (1) a member name or an alias of a module in the call library, and (2) it must be defined as an external name in the external symbol dictionary of the module with that name. If the unresolved external reference is a member name or an alias in the library, but is not an external name in that member, the member is processed but the external reference remains unresolved unless subsequently defined.

The automatic library-call mechanism searches the call library defined on the SYSLIB DD statement. The call library can contain either (1) object modules and control statements or (2) load modules; it must not contain both.

Modules from libraries other than the SYSLIB call library can be searched by the automatic library-call mechanism as directed by the LIBRARY control statement. The library specified in the control statement is searched for member names that match specific external references that are unresolved at the end of input processing. If any unresolved references are found in the modules located by automatic library call, they are resolved by another search of the library. Any external references not specified on a LIBRARY control statement are resolved from the library defined on the SYSLIB DD statement.

In addition, two means exist to negate the automatic library-call mechanism. The LIBRARY statement can be used to negate the automatic library call for selected external references unresolved after input processing; the NCAL option on the EXEC statement can be used to negate the automatic library

call for all external references unresolved after input processing. Use of the LIBRARY control statement and the NCAL option are discussed after the SYSLIB DD statement following.

SYSLIB DD STATEMENT

If the automatic library-call mechanism is to be used, the call library must be a partitioned data set described by a DD statement with a ddname of SYSLIB. Details concerning DCB requirements and record formats for SYSLIB libraries are given in "SYSLIB DD Statement" on page 57. The call library may be either a system call library or a private call library; call libraries may be concatenated.

System Call Library

For an example of some of the system programs that have their own automatic call library, see Figure 10. This library must be defined when an object module produced by that assembler or compiler is to be link-edited.

Program	Library Name
ALGOL	SYS1.ALGLIB
COBOL	SYS1.COBLIB
FORTTRAN	SYS1.FORTLIB
PL/I	SYS1.PL1LIB
Sort/Merge	SYS1.SORTLIB

Figure 10. System Automatic Call Libraries

The call library may contain input/output, data conversion, and/or other special routines (such as Sort/Merge SYS1.SORTLIB) that are needed to complete the module. The assembler or compiler creates an external reference for these special routines and the linkage editor resolves the references from the appropriate call library.

In the following example, a FORTRAN object module created in STEPA is to be link-edited in STEPB, and the FORTRAN automatic call library is used to resolve external references:

STEPA:		
//SYSOBJ	DD	DSNAME=&&OBJMOD,DISP=(NEW,
//		PASS),...
	.	
	:	
	.	
STEPB:		
//SYSLIN	DD	DSNAME=&&OBJMOD,DISP=(OLD,DELETE)
//SYSLIB	DD	DSNAME=SYS1.FORTLIB,DISP=SHR

The disposition of SHR on the SYSLIB DD statement means that other tasks that may be executing concurrently with STEPB may also use SYS1.FORTLIB.

Private Call Libraries

The SYSLIB DD statement can also describe a private, user-written library. In this case, the automatic library-call mechanism searches the private library for unresolved external references. In the following example, unresolved external references are to be resolved from a private library named PVTPROG:

```
//SYSLIB DD DSNAME=PVTPROG,DISP=SHR,UNIT=3380,  
// VOLUME=SER=PVT002
```

Concatenation of Call Libraries

System call libraries and private call libraries may be concatenated either to themselves, and/or to each other. When libraries are concatenated, they must all be either object module libraries or load module libraries; they may not be mixed.

If object modules from different system processors are to be link-edited to form one load module, the call library for each must be defined. This is accomplished by concatenating the additional call libraries to the library defined on the SYSLIB DD statement. In the following example, a FORTRAN object module and a COBOL object module are to be link-edited; the two system call libraries are concatenated as follows:

```
//SYSLIB DD DSNAME=SYS1.FORTLIB,DISP=SHR  
// DD DSNAME=SYS1.COBLIB,DISP=SHR
```

System libraries are cataloged; no unit or volume information is needed.

A system call library and a private call library can also be concatenated in this way. For example, by adding the following statement to the two in the preceding example, the private call library PVTPROG, which is not cataloged, is concatenated to the two system call libraries:

```
// DD DSNAME=PVTPROG,DISP=SHR,UNIT=3380,  
// VOLUME=SER=PVT002
```

Any external references not resolved from the two system libraries are resolved from the private library.

LIBRARY CONTROL STATEMENT

The LIBRARY control statement can be used to direct the automatic library-call mechanism to a library other than that specified in the SYSLIB DD statement. Only external references listed on the LIBRARY statement are resolved in this way. All other unresolved external references are resolved from the library in the SYSLIB DD statement.

The LIBRARY statement can also be used to specify external references that are not to be resolved by the automatic library-call mechanism. The LIBRARY statement specifies the duration of the nonresolution: either during the current linkage editor job step, called restricted no-call; or during this or any subsequent linkage editor job step, called never-call.

Examples of each use of the LIBRARY statement follow; a description of the format is given in "LIBRARY Statement" on page 79.

Additional Call Libraries

If the additional libraries are to be used to resolve specific references, the LIBRARY statement contains the ddname of a DD statement that describes the library. The LIBRARY statement also contains, in parentheses, the external references to be resolved from the library; that is, the names of the members to be used from the library. If the unresolved external reference is not a member name in the specified library, the reference remains unresolved unless subsequently defined.

For example, two modules (DATE and TIME) from a system call library have been rewritten. The new modules are to be tested with the calling modules before they replace the old modules. Because the automatic library call mechanism would otherwise search the system call library (which is needed for other modules), a LIBRARY statement is used, as follows:

```
//SYSLIB      DD      DSNAME=SYS1.COBLIB,DISP=SHR
//TESTLIB     DD      DSNAME=TEST,DISP=(OLD,KEEP),...
//SYSLIN      DD      DSNAME=ACCTROUT,...
//           DD      *
// LIBRARY    DD      TESTLIB(DATE,TIME)
/*
```

Two external references, DATE and TIME, are resolved from the library described on the TESTLIB DD statement. All other unresolved external references are resolved from the library described on the SYSLIB DD statement.

Restricted No-Call Function

The programmer can use the LIBRARY statement to specify those external references in the output module for which there is to be no library search during the current linkage editor job step. This is done by specifying the external reference(s) in parentheses without specifying a ddname. The reference remains unresolved, but the linkage editor marks the module executable.

For example, a program contains references to two large modules that are called from the automatic call library. One of the modules has been tested and corrected; the other is to be tested in this job step. Rather than execute the tested module again, the restricted no-call function is used to prevent automatic library call from processing the module as follows:

```

//          EXEC   PGM=HEWL,PARM=LET
//SYSLIB    DD      DSNAME=PVTPROG,DISP=SHR,UNIT=3380,
//          .      VOLUME=SER=PVT002
          .
//SYSLIN    DD      DSNAME=&&PAYROL,...
//          DD      *
LIBRARY     (OVERTIME)
/*

```

As a result, the external reference to OVERTIME is not resolved by automatic library call.

Never-Call Function

The never-call function specifies those external references that are not to be resolved by automatic library call during this or any subsequent linkage editor job step. This is done by specifying an asterisk followed by the external reference(s) in parentheses. The reference remains unresolved but the linkage editor marks the module executable.

For example, a certain part of a program is never executed, but it contains an external reference to a large module (CITYTAX) which is no longer used by this program. However, the module is in a call library needed to resolve other references. Rather than take up storage for a module that is never used, the never-call function is specified, as follows:

```

//          EXEC   PGM=HEWL,PARM=LET
//SYSLIB    DD      DSNAME=PVTPROG,DISP=SHR,UNIT=3380,
//          .      VOLUME=SER=PVT002
          .
//SYSLIN    DD      DSNAME=TAXROUT,DISP=OLD,...
//          DD      *
LIBRARY     *(CITYTAX)
/*

```

As a result, when program TAXROUT is link-edited, the external reference to CITYTAX is not resolved by automatic library call.

NCAL OPTION

When the NCAL option is specified, no automatic library call occurs to resolve external references that are unresolved after input processing. The NCAL option is similar to the restricted no-call function on the LIBRARY statement, except that the NCAL option negates automatic library call for all unresolved external references and restricted no-call negates automatic library call for selected unresolved external references. With NCAL, all external references that are unresolved after input processing is finished, remain unresolved. The module is, however, marked executable.

The NCAL option is a special processing parameter that is specified on the EXEC statement as described in "No Automatic Library-Call Option" on page 45.

INCLUDED DATA SETS

The INCLUDE control statement requests the linkage editor to use additional data sets as input. These can be sequential data sets containing object modules and/or control statements, or members of partitioned data sets containing object modules and/or control statements, or load modules.

The INCLUDE statement specifies the ddname of a DD statement that describes the data set to be used as additional input. If the DD statement describes a partitioned data set, the INCLUDE statement also contains the name of each member to be used. See "INCLUDE Statement" on page 76 for a detailed description of the format of the INCLUDE statement.

When an INCLUDE control statement is encountered, the linkage editor processes the module or modules indicated. Figure 11 shows the processing of an INCLUDE statement. In the illustration, the primary input data set is a sequential data set named OBJMOD which contains an INCLUDE statement. After processing the included data set, the linkage editor processes the next primary input item. The arrows indicate the flow of processing.

If an included data set also contains an INCLUDE statement, this specified module is also processed. However, any data following the INCLUDE statement is not processed.

If the OBJMOD data set shown in Figure 11 is itself included, the data following the INCLUDE statement for OBJLIB is not processed. Figure 12 on page 34 shows the flow of processing for this example.




Figure 11. Processing of One INCLUDE Control Statement




Figure 12. Processing of More than One INCLUDE Control Statement

Including Sequential Data Sets

Sequential data sets containing object modules and/or control statements can be specified by an INCLUDE control statement. In the following example, an INCLUDE statement specifies the ddnames of two sequential data sets to be used as additional input:

```

//ACCOUNTS DD DSNAME=ACCTROUT,DISP=(OLD,KEEP),...
//INVENTORY DD DSNAME=INVENTORY,DISP=(OLD,KEEP),...
//SYSLIN DD DSNAME=QTREND,...
// DD *
INCLUDE ACCOUNTS,INVENTORY
/*

```

Each ddname could also have been specified on a separate INCLUDE statement; with either method, a DD statement must be specified for each ddname.

Another method of doing the preceding example is given in "Including Concatenated Data Sets" on page 35.

Including Library Members

One or more members of a partitioned data set can be specified on an INCLUDE control statement. The member name must be specified on the INCLUDE statement; no member name should appear on the DD statement itself.

In the following example, one member name is specified on the INCLUDE statement:

```
//PAYROLL DD DSN=PAYROUTS,DISP=(OLD,KEEP),...
//SYSLIN DD DSN=CHECKS,DISP=(OLD,DELETE),...
// DD *
INCLUDE PAYROLL(FICA)
/*
```

If more than one member of a partitioned data set is to be included, the INCLUDE statement specifies all the members to be used from each library. The member names appear in parentheses, following the data set name of the library. The member names are not repeated on the DD statement.

In the following example, an INCLUDE statement specifies two members from each of two libraries to be used as additional input:

```
//PAYROLL DD DSN=PAYROUTS,DISP=(OLD,KEEP),...
//ATTEND DD DSN=ATTROUTS,DISP=(OLD,KEEP),...
//SYSLIN DD *
INCLUDE PAYROLL(FICA,TAX),ATTEND(ABSENCE,OVERTIME)
/*
```

Each library could have been specified on a separate INCLUDE statement; with either method, a DD statement must be specified for each ddname.

Another method of doing this example is given in "Including Concatenated Data Sets."

Including Concatenated Data Sets

Several data sets can be designated as input with one INCLUDE statement that specifies one ddname; additional data sets are then concatenated to the data set described on the specified DD statement. When data sets are concatenated, all records must have the same characteristics (that is, format, record length, block size, and so forth).

SEQUENTIAL DATA SETS: In the following example, two sequential data sets are concatenated and then specified as input with one INCLUDE statement:

```
//CONCAT DD DSN=ACCTROUT,DISP=(OLD,KEEP),...
// DD DSN=INVENTORY,DISP=(OLD,KEEP),...
//SYSLIN DD DSN=SALES,DISP=OLD,...
// DD *
INCLUDE CONCAT
/*
```

When the INCLUDE statement is recognized, the contents of the sequential data sets ACCTROUT and INVENTORY are processed.

LIBRARY MEMBERS: Members from more than one library can be designated as input with one ddname on an INCLUDE statement. In this case, all the members are listed on the INCLUDE statement; the partitioned data sets are concatenated using the ddname from the INCLUDE statement:

```
//CONCAT DD DSNAME=PAYROUTS,DISP=(OLD,KEEP),...
// DD DSNAME=ATTROUTS,DISP=(OLD,KEEP),...
//SYSLIN DD DSNAME=REPORT,DISP=OLD,...
// DD *
INCLUDE CONCAT(FICA,TAX,ABSENCE,OVERTIME)
/*
```

When the INCLUDE statement is recognized, the two libraries, PAYROUTS and ATTROUTS, are searched for the four members; the members are then processed as input.

CHAPTER 4. SPECIFYING JCL TO RUN A LINKAGE EDITOR JOB

This chapter summarizes those aspects of the job control language that pertain directly to the use of the linkage editor. The major topics covered are the EXEC statement, DD statements, and cataloged procedures for the linkage editor. The reader should be familiar with the job control language as described in the publication JCL.

EXEC STATEMENT—INTRODUCTION

The EXEC statement is the first statement of every job step. For the linkage editor job step, the following topics are pertinent:

- The program name of the linkage editor
- Linkage editor options passed to the job step
- Region-size requirements for the linkage editor

For an execution job step following the linkage editor job step, the linkage editor return code is important.

The EXEC statement contains the symbolic name of the load module to be invoked for execution. The linkage editor can be invoked with the following program name:

HEWL

LINKEDIT is an alias name for the linkage editor and can also be used to invoke it.

For example, the following EXEC statement causes the linkage editor to be invoked:

```
//LKED      EXEC  PGM=HEWL
```

PGM=LINKEDIT could also be used.

To ensure compatibility with the operating system, the linkage editor can also be invoked by any of the following alias names: IEWL, IEWLF440, IEWLF880, and IEWLF128.

EXEC STATEMENT—JOB STEP OPTIONS

The EXEC statement also contains a list of options or parameters to be passed to the linkage editor. These options are of four types:

- Module attributes, which describe the characteristics of the output load module
- Special processing options, which affect linkage editor processing
- Space allocation options, which affect the amount of storage used by the linkage editor for processing and output module library buffers
- Output options, which specify the kind of output the linkage editor is to produce

The rest of this section describes the options in each category. All the options for a particular linkage editor execution are listed in the PARM parameter on the EXEC statement. They can be listed in any sequence, as long as the rules for coding parameters are followed.

MODULE ATTRIBUTES

The module attributes describe the characteristics of the output module, or modules. (If more than one load module is produced by the same linkage editor job step, all output modules will have the attributes assigned on the EXEC statement.) The attributes for each load module are stored in the directory of the output module library along with the member name. (The format of the directory entry of a partitioned data set is given in Data Areas—JES3.)

Module attributes specify whether or not the module:

- Can ever be processed by the linkage editor
- Can be brought into virtual storage only by the LOAD macro instruction
- Is to be in overlay format
- Can be reused
- Can be placed in the link pack area; that is, is reenterable
- Can be replaced during execution by recovery management; that is, is refreshable
- Is to be tested by the TSO TEST command
- Is to have specified control sections aligned on page boundaries
- Is or is not authorized to use the restricted system resources and functions

After the descriptions of the module attributes, the default and incompatible attributes are discussed.

Downward Compatible Attribute

When this attribute is specified, a maximum record size of 1024 bytes is used for the output module library.

To assign the downward compatible attribute, code DC in the PARM field as follows:

```
//LKED EXEC PGM=IEWL,PARM='DC,...'
```

Notes:

If the DC attribute is specified and the output load module library is a data set created by the link-edit job step, the blocksize in the DSCB (data set control block) is set to 1024. If the DC attribute is specified and the output load module library is an existing data set, then the blocksize in the DSCB is set to 1024 only if the current blocksize in the DSCB is less than 1024; if the current blocksize in the DSCB is greater than 1024, the load module is written using a maximum record size of 1024 bytes but the blocksize in the DSCB is not changed.

Scatter Format Attribute

When the scatter format attribute is specified, the linkage editor produces a load module in a format suitable for either scatter or block loading.

To assign the scatter format attribute, code SCTR in the PARM field, as follows:

```
//LKED EXEC PGM=IEWL,PARM='SCTR,...'
```

Notes:

1. If scatter format is not specified, the block format attribute is assigned by the linkage editor. (The programmer cannot specify block format.)
2. If SCTR is specified, the programmer should ensure that the load module does not contain zero-length control sections, private code sections, or common areas. The presence of such sections in a module that is to be scatter loaded can, under certain circumstances, cause the module to be loaded incorrectly.
3. The SCTR attribute must be specified when the nucleus for a VS system is link-edited. In all other instances, if the SCTR attribute is specified, the linkage editor builds the output load module appropriately; however, scatter load support is not provided in VS systems and the attribute/load module format is ignored when fetching the load module.

Not Editable Attribute

A load module which is marked NE (not editable) is not reprocessible by the linkage editor. If a module map or a cross-reference table is requested, the not-editable attribute is ignored.

To assign the not-editable attribute, code NE in the PARM field, as follows:

```
//LKED      EXEC   PGM=HEWL,PARM='NE,...'
```

Note: The not-editable attribute disables the EXPAND function for the output load module and also limits to 18 the number of consecutive iterations of AMASPZAP. If the EXPAND function is required or more than 18 iterations of AMASPZAP are required, the load module must be re-created.

Only-Loadable Attribute

A module with the only-loadable attribute can be brought into virtual storage only with a LOAD macro instruction. Some subsets of the control program use a smaller control table when the load module is invoked with a LOAD. This reduces the overall virtual storage requirements of the module.

The module with the only-loadable attribute must be entered by means of a branch instruction or a CALL macro instruction. If an attempt is made to enter the module with a LINK, XCTL, or ATTACH macro instruction, the program making the attempt is terminated abnormally by the control program.

To assign the only-loadable attribute, code OL in the PARM field as follows:

```
//LKED      EXEC   PGM=HEWL,PARM='OL,...'
```

Overlay Attribute

A program with the overlay attribute is placed in an overlay structure as directed by linkage editor OVERLAY control statements. The module is suitable only for block loading; it cannot be refreshable, reenterable, or serially reusable.

If the overlay attribute is specified and no OVERLAY control statements are found in the linkage editor input, the attribute is negated. The condition is considered a recoverable error; that is, if the LET option is specified, the module is marked executable.

The overlay attribute must be specified for overlay processing. If this attribute is omitted, the OVERLAY and INSERT statements are considered invalid, and the module is not an overlay structure. This condition is also recoverable; if the LET option is specified, the module is marked executable.

To assign the overlay attribute, code OVLY in the PARM field as follows:

```
//LKED      EXEC   PGM=HEWL,PARM='OVLY,...'
```

See "Appendix C. Designing and Specifying Overlay Programs" on page 139, for information on the design and specification of an overlay structure.

Reusability Attributes

Either one of two attributes may be specified to denote the reusability of a module. (Reusability means that the same copy of a load module can be used by more than one task either concurrently or one at a time.) The reusability attributes are reenterable and serially reusable; if neither is specified, the module is not reusable and a fresh copy must be brought into virtual storage before another task can use the module.

The linkage editor only stores the attribute in the directory entry; it does not check whether the module is really reenterable or serially reusable. A reenterable module is automatically assigned the reusable attribute. However, a reusable module is not also defined as reenterable; it is reusable only.

REENTERABLE: A module with the reenterable attribute can be executed by more than one task at a time; that is, a task may begin executing a reenterable module before a previous task has finished executing it. This type of module cannot be modified by itself or by any other module during execution.

If a module is to be reenterable, all the control sections within the module must be reenterable. If the reenterable attribute is specified, and any load modules that are not reenterable become a part of the input to the linkage editor, the attribute is negated.

To assign the reenterable attribute, code RENT in the PARM field, as follows:

```
//LKED      EXEC   PGM=HEWL,PARM='RENT,...'
```

SERIALY REUSABLE: A module with the serially reusable attribute can be executed by only one task at a time; that is, a task may not begin executing a serially reusable module before a previous task has finished executing it. This type of module must initialize itself and/or restore any instructions or data in the module altered during execution.

If a module is to be serially reusable, all its control sections must be either serially reusable or reenterable. If the serially reusable attribute is specified, and any load modules that are neither serially reusable nor reenterable become a part of the input to the linkage editor, the serially reusable attribute is negated.

To assign the serially reusable attribute, code REUS in the PARM field, as follows:

```
//LKED      EXEC   PGM=HEWL,PARM='REUS,...'
```

Refreshable Attribute

A module with the refreshable attribute can be replaced by a new copy during execution by a recovery management routine without changing either the sequence or results of processing. This type of module cannot be modified by itself or by any other module during execution. The linkage editor only stores the attribute in the directory entry; it does not check whether the module is refreshable.

If a module is to be refreshable, all the control sections within it must be refreshable. If the refreshable attribute is specified, and any load modules that are not refreshable become a part of the input to the linkage editor, the attribute is negated.

To assign the refreshable attribute, code REFR in the PARM field, as follows:

```
//LKED      EXEC   PGM=HEWL,PARM='REFR,...'
```

Test Attribute

A module with the test attribute is to be tested and contains the testing symbol tables for the TSO TEST command. The linkage editor accepts these tables as input, and places them in the output module. The module is marked as being under test. If the test attribute is not specified, the symbol tables are ignored by the linkage editor and are not placed in the output module. If the test attribute is specified, and no symbol table input is received, the output load module will not contain symbol tables to be used by the TSO TEST command.

To assign the test attribute, code TEST in the PARM field, as follows:

```
//LKED      EXEC   PGM=HEWL,PARM='TEST,...'
```

Note: The test attribute applies to programs using either TESTRAN or the TSO TEST command. Do not use the 'TEST' option unless the load module is to be executed by either TSO or TESTRAN.

Authorization Code

The output load module is assigned an authorization code that determines whether or not the load module may use restricted system services and resources.

To assign an authorization code through the PARM field, code the AC parameter as follows:

```
//LKED      EXEC   PGM=HEWL,PARM='AC=n,...'
```

The authorization code, n, must be 1 to 3 decimal digits with a value from 0 to 255.

'AC=,...' and 'AC= ' are equivalent to 'AC=0'. The authorization code assigned in the PARM field is overridden by an authorization code assigned through the SETCODE control statement.

Addressing Mode Attribute

To assign the addressing mode for all the entry points into the load module (the main entry point, its true aliases, and all the alternate entry points), code the AMODE parameter as follows:

```
//LKED EXEC PGM=IEWL,  
          PARM='AMODE=xxx,...'
```

The addressing mode 'xxx' must be either 24, 31, or ANY.

The addressing mode assigned in the PARM field overrides the separate addressing modes found in the ESD data for the control sections or private code where the entry points are located. The addressing mode assigned in the PARM field is overridden by an addressing mode assigned in the MODE control statement.

If the AMODE parameter occurs more than once in the PARM field of the EXEC statement, the last valid parameter is used.

If only the AMODE value is specified in the PARM field of the EXEC statement, an RMODE value of 24 is implied.

Note: The keyword 'AMODE' may also be specified as 'AMOD'.

Residence Mode Attribute

To assign the residence mode for the output load module, code the RMODE parameter as follows:

```
//LKED EXEC PGM=IEWL,  
          PARM='RMODE=xxx,...'
```

The residence mode 'xxx' must be either 24 or ANY.

The residence mode assigned in the PARM field overrides the residence mode accumulated from the input control sections and private code. The residence mode assigned in the PARM field is overridden by a residence mode assigned through the MODE control statement.

If the RMODE parameter occurs more than once in the PARM field of the EXEC statement, the last valid parameter is used.

If only an RMODE value of ANY is specified in the PARM field of the EXEC statement, an AMODE value of 31 is implied.

If only an RMODE of 24 is specified, no overriding AMODE value is assigned; instead, the AMODE value in the ESD data for the main entry point, a true alias, or an alternate entry point is used in generating its respective directory entry. If any control section to be linked has an RMODE=24, then the load module is marked RMODE=24.

Note: The keyword 'RMODE' may also be specified as 'RMODE'.

AMODE/RMODE Combinations in the PARM Field

In generating a directory entry for the main entry point, a true alias, or an alternate entry point, the linkage editor validates the combination of the AMODE value and the RMODE value, as specified by the user in the PARM field of the EXEC statement, according to the following table:

	RMODE=24	RMODE=ANY
AMODE=24	valid	invalid
AMODE=31	valid	valid
AMODE=ANY	valid	invalid

If the AMODE/RMODE combination resulting from the PARM field of the EXEC statement is invalid, an error message is issued and the linkage editor ignores the PARM field of the EXEC statement as the source of AMODE/RMODE data.

Default Attributes

Unless specific module attributes are indicated by the programmer, the output module is not in an overlay structure, and it is not tested. The module is in block format, not refreshable, not reenterable, and not serially reusable. If page boundary alignment is requested, its control sections are aligned on 4K-byte page boundaries.

One other attribute is specified by the linkage editor after processing is finished. If, during processing, severity 2 errors were found that would prevent the output module from being executed successfully, the linkage editor assigns the not-executable attribute. The control program will not load a module with this attribute.

If the LET option is specified, the output module is marked executable even if severity 2 errors occur. (The LET option is discussed later in this section.)

If the AC parameter is not specified or is coded incorrectly, the default authorization code of 0 is assigned to the output load module.

Incompatible Attributes

Of the module attributes the programmer may specify, several are mutually exclusive. When mutually exclusive attributes are specified for a load module, the linkage editor ignores the less-significant attributes. For example, if both OVLY and RENT are specified, the module will be in an overlay structure and will not be reenterable.

Certain attributes are also incompatible with other job step options. All job step options are shown in Figure 15 on page 53 along with those options that are incompatible.

SPECIAL PROCESSING OPTIONS

The special processing options affect the ability to execute the output module and the use of the automatic library-call mechanism. These options are the exclusive call option, the let execute option, and the no automatic-call option.

Exclusive Call Option

When the exclusive call option is specified, valid exclusive references have been made between segments, and the linkage editor marks the output module as executable. However, a warning message is given for each valid exclusive reference.

To specify the exclusive call option, code XCAL in the PARM field as follows:

```
//LKED EXEC PGM=HEWL,PARM='XCAL,OVLY,...'
```

The OVLY attribute must also be specified for an overlay program.

Note: Unless the let execute option is specified, other errors may cause the module to be marked not executable.

Let Execute Option

When the let execute option is specified, the linkage editor marks the output module as executable even though a severity 2 error condition was found during processing. (A severity 2 error condition could make execution of the output load module impossible.) Some examples of severity 2 errors are:

- Unresolved external references
- Valid or invalid exclusive calls in an overlay program
- Error on a linkage editor control statement
- A library module that cannot be found
- No available space in the directory of the output module library

To specify the let execute option, code LET in the PARM field as follows:

```
//LKED      EXEC   PGM=HEWL,PARM='LET,...'
```

Note: If LET is specified, XCAL need not be specified.

No Automatic Library-Call Option

When the no automatic library-call option is specified, the linkage editor library-call mechanism does not call library members to resolve external references. The output module is marked executable even though unresolved external references are present. If this option is specified, the LIBRARY statement need not be used to negate the automatic library call for selected external references. Also, with this option, a SYSLIB DD statement need not be supplied.

To specify the no automatic library-call option, code NCAL in the PARM field, as follows:

```
//LKED      EXEC   PGM=HEWL,PARM='NCAL,...'
```

Note: Unless the LET option is also specified, other errors may cause the module to be marked not executable.

SPACE ALLOCATION OPTIONS

These options allow the programmer to specify the storage available to the linkage editor, and to specify the block size for the output module. For large modules and SMP, see SMP System Programmer's Guide; for SMP/E, see SMP/E User's Guide.

SIZE Option

The programmer can specify, through the SIZE option, the amount of virtual storage to be used by the linkage editor and the portion of that storage to be used as the load module buffer.

The linkage editor provides default values for the SIZE option. The default values are used if one or both of the values are not specified correctly by the user or are not specified at all. These defaults should be adequate for most link-edits, relieving the programmer from specifying the SIZE option for each link-edit. The default values are: value1 is 384K bytes and value2 is 96K bytes.

FORMAT: The format of the SIZE option is:

```
SIZE=(value1,value2)
```

```
SIZE=(value1)
```

```
SIZE=(value1,)
```

```
SIZE=(,value2)
```

```
SIZE=(,)
```

When coded in the PARM field, value1 and value2 parameters are enclosed in parentheses as follows:

```
//LKED      EXEC   PGM=HEWL,  
//          PARM='SIZE=(value1,value2),...'
```

Both value1 and value2 may be expressed as integers specifying the number of bytes of virtual storage or as nK, where n represents the number of 1K (1024) bytes of virtual storage.

When determining the values for the SIZE option, it is best to establish value2 first, then value1.

VALUE2: Value2 specifies the number of bytes of storage to be allocated as the load module buffer. The allocation specified by value2 is a part of the virtual storage specified by value1.

The actual minimum for value2 is 6144 (6K) or the length of the largest input load module text record, whichever is larger. AMBLIST may be used to find the size of the load module text records. If a value less than 6144 (6K) is specified, the default value of 96K for value2 is used.

The space allocated by value2 is used for: the buffer into which the input load module text is read, the buffer from which load module text is written to the intermediate data set, the buffer into which the load module text is read from the intermediate data set, and the buffers from which the load module text is written to the output data set. Therefore, the determination of value2 requires that the programmer consider the record sizes of the data sets from which any load module text records are to be read (SYSLIB, any data set referenced by an INCLUDE, any library data set), the record size for the intermediate data set (SYSUT1), and the record size for the output load module data set (SYSLMOD).

Figure 13 lists the direct access devices that may contain data sets that are the source of input load module text, the intermediate data set, and the output load module data set, and lists the maximum record size used for each device by the linkage editor. These maximum record sizes may always be used in specifying value2 or, if the programmer can determine them, exact sizes can be used.

Device	Device Maximum Record Size (Bytes)	SYSUT1 or SYSLMOD Maximum Record Size (Bytes)
2305-2	14660	13312
3330-1	13030	12288
3330-11	13030	12288
3340	8368	7680
3344	8368	7680
3350	19069	18432
3375	32760	18432
3380 ¹	32760	18432

Figure 13. SYSUT1 and SYSLMOD Device Types and Their Maximum Record Sizes

Note to Figure 13:

¹ 3380 models A04, AA4, B04, AD4, BD4, AE4, and BE4.

The programmer must specify value2 so that the linkage editor has sufficient space to allocate buffers that are compatible with the record sizes for the intermediate data set and the output load module data set.

The linkage editor optimizes the record size for the device type of output load module data set unless one of the following conditions exists.

1. The programmer has specified PARM='...DC,...', forcing the linkage editor to write records having a maximum size of 1024 (1K) bytes.
2. The programmer has specified PARM='...DCBS,...', and the SYSLMOD DD statement contains a BLKSIZE subparameter in the DCB parameter, forcing the linkage editor to write records having a maximum length equal to the BLKSIZE specification.
3. The output load module data set is an existing data set having a block size less than the optimum record size, forcing the linkage editor to write records no longer than that block size.
4. The programmer has specified a value2 less than twice the maximum record size for the output load module data set, forcing the linkage editor to write records having a maximum size of one-half value2.
5. The intermediate data set and the output load module data set have dissimilar record sizes, forcing the linkage editor to write records having a maximum size determined for compatibility between the two data sets.

The linkage editor optimizes the record size of the output load module data set for its device type but selects a record size compatible with the intermediate data set (see restrictions above). Therefore, if the intermediate data set and the output load module data set reside on the same device type, use of the load module buffer is optimized. Also, if the data sets are on different units of the same type, the performance of the linkage editor is improved.

Figure 14 shows the record sizes used for compatibility between every combination of device types for the intermediate and output load module data sets.

SYSLMOD Record Size		SYSUT1 Record Size		
Device Used	Maximum Record Size Produced	Device Used	Maximum Record Size Produced	Minimum Load Module Buffer Area (Value2)
IBM 2305-2	13K	2305-2	13K	26K
	12K ¹	3330, 3330-11	12K	24K
	7.5K ¹	3340	7.5K	15K
	13K	3350, 3375	13K ²	26K
	13K	3380 ³	13K ²	26K
IBM 3330	12K	2305-2	12K ²	24K
IBM 3330-11	12K	3330, 3330-11	12K	24K
	7.5 ¹	3340	7.5K	15K
	12K	3350, 3375	12K ²	24K
	12K	3380 ³	12K ²	24K
IBM 3340	7.5K	2305-2	7.5K ²	15K
IBM 3344	7.5K	3330, 3330-11	7.5K ²	15K
	7.5K	3340	7.5K	15K
	7.5K	3350, 3375	7.5K ²	15K
	7.5K	3380 ³	7.5K ²	15K
IBM 3350	13K ¹	2305-2	13K	26K
IBM 3375	12K ¹	3330, 3330-11	12K	24K
IBM 3380 ³	15K ¹	3340	7.5K	30K
	18K	3350, 3375	18K	36K
	18K	3380 ³	18K	36K

Figure 14. Load Module Buffer Area and SYSLMOD and SYSUT1 Record Sizes

Notes to Figure 14:

- ¹ The SYSLMOD record size is reduced to less than the maximum to make it compatible with the SYSUT1 record size.
- ² The SYSUT1 record size is reduced to less than the maximum to make it compatible with the SYSLMOD record size.
- ³ 3380 models A04, AA4, B04, AD4, BD4, AE4, and BE4.

Value2 should be, minimally, twice the record size for the output load module data set. If value2 can be made larger than twice the record size for the output load module data set, the increase should be the larger of the record sizes for the intermediate and output load module data sets.

The practical maximum for value2 is the length of the load module to be built, plus 4K bytes if the length of the load module to be built is equal to or greater than 40960 (40K). Any space allocated to the load module buffer above this amount is not used and does not need be allocated to value2.

If a value2 is specified that cannot be accommodated in the available storage, value2 is reduced to the next lower 2K-byte multiple of storage that is available. This reduction, however, never decreases value2 to less than the minimum, 6144 (6K).

The optimal value2 is the practical maximum, as explained above. If the entire load module is contained in storage, the performance of the linkage editor is improved and the use of the intermediate data set may be eliminated.

Examples of Value2 Determination

1. A load module of between 21K and 22K bytes is to be built. The load module data set is a new data set on an IBM 3330 Disk Storage device. The intermediate data set is allocated to an IBM 3340 Direct Access Storage device. A SYSLIB data set is to be used, residing on a 3330. The entire load module could be contained in the load module buffer if value2 were 22K bytes (the load module size). The practical minimum for value2 would be 12K bytes (the size of the largest possible input load module text record from the SYSLIB data set). However, value2 should be at least as large as two records to be written to the load module data set (that is, 24K bytes). There is a reconciliation necessary in this case between the two dissimilar device types for the intermediate and output load module data sets; but the record size of the output load module data set is an even multiple of the record size of the intermediate data set so no adjustment of the record sizes is made. Therefore, the practical minimum, as well as the practical maximum and optimal value2 in this case is 24K bytes.
2. A load module of more than 50K bytes is to be relink-edited; however, a maximum of 40K bytes is available to be allocated to value2. The output load module data set is an old data set residing on a 3340, written with maximum record size. The intermediate data set is allocated to an IBM 2305-2 Fixed Head Storage device. The link-edit involves a control section in the SYSLIN data set that will replace a control section in the old load module, followed by an INCLUDE statement naming the old load module on the SYSLMOD data set. The maximum for value2 cannot be satisfied, since only 40K bytes is available. The size of two maximum records written to a 3340 would be 14K bytes. However, the size of one record to be written or to be read from the intermediate data set is 14K bytes. Therefore, the minimum for value2 in this case is 14K bytes. This is sufficient space for one input load module text record or one record written to or to be read from the intermediate data set or two records written to the output load module data set.

3. The output load module data set resides on a 2305-2. The intermediate data set is allocated to a 3330. All load module input comes from a 3330. Value2 in this case is 24K bytes, because the input load module text records are, at most, 12K bytes, the records written to and read from the intermediate data set are 12K bytes, and the records written to the output load module data set are 12K bytes. The maximum record size of 14K bytes for the 2305-2 is reduced to 12K bytes for this link-edit in order to be compatible with the intermediate data set.

An alternative for value2 in the above example is 12K bytes. This 12K bytes is adequate for the input load module text records and the records written to and read from the intermediate data set. The 12K value forces a maximum record size of 6K bytes to be written to the output load module data set. At 6K bytes each, two records can be written on a 2305-2 track while, as in the above example, only one record of 12K bytes can be written on a 2305-2 track.

4. A load module of 10K is to be link-edited. The output load module data set resides on a 2305. The input load module libraries all reside on 2314s. The intermediate data set is allocated to a 2314. The programmer has specified the linkage editor parameter DC. The minimum for value2 of 6K is adequate in this case, since 6K is sufficient for input and intermediate data set records and the output load module data sets records have a maximum size of 1K.
5. The output load module data set is a new data set allocated to a 3330. The programmer has specified the linkage editor parameter DCBS, and the SYSLMOD DD statement contains '...DCB=...BLKSIZE=3072,......'. The only load module input comes from a data set created previously in a similar manner. The intermediate data set is allocated to a 3340. The minimum for value2 in this case is 6K bytes; the input load module records are 3K bytes at most, the intermediate data set records are 7K bytes at most, and, as directed by the programmer, the linkage editor produces records having a maximum size of 3K bytes on the output load module data set.

VALUE1: Value1 specifies the number of bytes of virtual storage available to the linkage editor regardless of the private area size. The storage specified by value1 includes the allocation specified by value2.

The absolute minimum for value1 is the design point of the linkage editor, 96K bytes. If a value less than the minimum for value1 is specified, the default options for both value1 and value2 are used.

The practical minimum for value1 is 98304 (96K) bytes plus any excess in value2 over 6144 (6K) bytes, plus any additional space required to support the blocking factor for the SYSLIN, object module library, and SYSPRINT data sets.

The design point of the linkage editor provides for the minimum load module buffer—6144 (6K) bytes of virtual storage. If a load module buffer larger than 6144 (6K) bytes is specified in value2, value1 must be increased by the excess of that value2 over 6144 (6K) bytes.

The linkage editor supports three different blocking factors for the SYSLIN, object module library, and SYSPRINT data sets; they are 5, 10, and 40 to 1. The requirement for additional space depends upon the blocking factor that is to be supported.

The following table shows the additional space required to support each blocking factor.

Blocking Factor	Space Required
5 to 1	0 or 0K
10 to 1	18432 or 18K
40 to 1	28672 or 28K

Blocking factors of 1 through 4, 6 through 9, and 11 through 39 are treated as blocking factors of 5, 10, and 40, respectively. Blocking factors greater than 40 are invalid.

The additional space requirement is determined by the largest blocking factor among the affected data sets.

The blocking factor supported is dependent upon space available after value2 has been allocated to the load module buffer out of value1. Therefore, if the space provided in value1 is insufficient, the next smallest blocking factor is used.

The performance of the linkage editor can be improved by the allocation of additional storage by value1, especially in providing for the optimal value2.

The maximum value that can be specified for value1 is 9999999 or 9999K. However, the amount of virtual storage actually allocated for value1 is the smaller of:

- The region size
- The amount specified for value1

Examples of Value1 Determination

1. Assume that an optimum value2 of 36K bytes has already been determined for the link-edit. An appropriate value1 is 126K bytes, because an additional 30K bytes, above the minimum of 96K bytes, is needed to support the allocation of 36K bytes to value2 and no additional storage is required to support the blocking factors for SYSLIN, SYSPRINT, and any object module libraries.
2. The minimum for value2 (6K bytes) is used. All the object module libraries are blocked 5-to-1, except one that is blocked 10-to-1. The SYSLIN and SYSPRINT data sets are assigned blocking factors of 5. An appropriate value1 for this link-edit is 114K bytes, the minimum plus the 18K bytes needed to support the blocking factor of 10-to-1 on the object module library.

DCBS Option

The DCBS option allows the programmer to specify the block size for the SYSLMOD data set in the DCB parameter of SYSLMOD DD statement.

If the DCBS option is specified, the block size value in the DSCB for the SYSLMOD data set may be overridden. If the DCBS option is not specified, the block size value in the DSCB for the SYSLMOD data set may not be overridden.

If the DCBS option is specified and no block size value is provided in the DCB parameter of the SYSLMOD DD statement, the linkage editor uses the maximum track size for the device. If the DCBS option is not specified and a block size value is provided in the DCB parameter of the SYSLMOD DD statement, the block size value in the DCB parameter of the SYSLMOD DD statement is ignored by the linkage editor.

Even though the DCBS option is specified, the linkage editor will not allow the programmer to set the block size for the SYSLMOD data set to a value less than the minimum; that is, 256, or 1024 if the SCTR option is specified, or a value less than the block size in the DSCB for an existing data set.

The block size specified by the programmer will be used unless (1) it is larger than the maximum record size for the device, in which case the maximum record size is used, or (2) it is less than the minimum block size, in which case the minimum block size is used.

The following example shows the use of the DCBS option for an IBM 3380 Direct Access Storage device:

```
//LKED      EXEC  PGM=HEWL,PARM='XREF,DCBS'  
           .  
           .  
//SYSLMOD   DD    DSNAME=LOADMOD(TEST),DISP=(NEW,KEEP),  
//          DCB=(BLKSIZE=23440),...
```

As a result, the linkage editor uses a 23440 block size for the output module library.

OUTPUT OPTIONS

These options control the optional diagnostic output produced by the linkage editor. The programmer can request that the linkage editor produce a list of all control statements and a module map or cross-reference table to help in testing a program. The format of each is described in "Chapter 8. Interpreting Linkage Editor Output" on page 109.

In addition, the programmer can request that the numbered error/warning messages generated by the linkage editor appear on the SYSTERM data set as well as on the SYSPRINT data set.

Control Statement Listing Option

To request a control statement listing, code LIST in the PARM field, as follows:

```
//LKED      EXEC  PGM=HEWL,PARM='LIST,...'
```

When the LIST option is specified, all control statements processed by the linkage editor are listed in card-image format on the diagnostic output data set.

Module Map Option

To request a module map, code MAP in the PARM field, as follows:

```
//LKED      EXEC  PGM=HEWL,PARM='MAP,...'
```

When the MAP option is specified, the linkage editor produces a module map of the output module on the diagnostic output data set.

Cross Reference Table Option

To request a cross-reference table, code XREF in the PARM field, as follows:

```
//LKED      EXEC   PGM=HEWL,PARM='XREF,...'
```

When the XREF option is specified, the linkage editor produces a cross-reference table of the output module on the diagnostic output data set. The cross-reference table includes a module map; therefore, both XREF and MAP need not be specified for one linkage editor job step.

Alternate Output (SYSTEM) Option

To request that the numbered linkage editor error/warning messages be generated on the data set defined by a SYSTEM DD statement, code TERM in the PARM field, as follows:

```
//LKED      EXEC   PGM=HEWL,PARM='TERM,...'
```

When the TERM option is specified, a SYSTEM DD statement must be provided. If it is not, the TERM option is negated.

Output specified by the TERM option supplements printed diagnostic information; when TERM is used, linkage editor error/warning messages appear in both output data sets.

INCOMPATIBLE JOB STEP OPTIONS

When mutually exclusive job step options are specified for a linkage editor execution, the linkage editor ignores the less significant options. Figure 15 on page 53 illustrates the significance of those options that are incompatible. When an X appears at an intersection, the options are incompatible. The option that appears higher in the list is selected.

Note: An X indicates incompatible attributes; the attribute that appears lower on the list is ignored. For example, to check the compatibility of XREF and NE, follow the XREF column down and the NE row across until they intersect. Because an X appears where they intersect, they are incompatible attributes. NE is ignored.

Figure 15. Incompatible Job Step Options for the Linkage Editor

For example, to check the compatibility of XREF and NE, follow the XREF column down and the NE row across until they intersect. Because an X appears where they intersect, they are incompatible; XREF is selected; NE is negated.

If incorrect values are specified for the SIZE parameter, the default values are used. If incompatible options are detected, the message

```
*** OPTIONS INCOMPATIBLE ***
```

is printed. This message follows the standard module disposition message.

If the incompatible options OVLY and AMODE or RMODE are specified, a diagnostic message is issued.

EXEC STATEMENT—REGION PARAMETER

The REGION parameter specifies the maximum amount of storage that can be allocated to satisfy a request for storage that the linkage editor makes. In its minimal situation, the linkage editor requires a REGION parameter of not less than 96K bytes; in its default situation, not less than 512K bytes; and, in its maximal situation (see "Size Parameter Guidelines" on page 61), not less than 1500K bytes.

EXEC STATEMENT—RETURN CODE

The linkage editor passes a return code to the control program upon completion of the job step. The return code reflects the highest severity code recorded in any iteration of the linkage editor within that job step. The highest severity code encountered during processing is multiplied by 4 to create the return code; this code is placed into register 15 at the end of linkage editor processing. Figure 16 contains the return codes, the corresponding severity code, and a description of each.

Return Code	Severity Code	Description
00	0	Normal conclusion
04	1	Warning messages have been listed; execution should be successful. For example, if the overlay option is specified and the overlay structure contains only one segment, a return code of 04 is placed in register 15.
08	2	Error messages have been listed; execution may fail. The module is marked not executable unless the LET option is specified. For example, if the block size of a specified library data set cannot be handled by the linkage editor, a return code of 08 is placed in register 15.
0C	3	Severe errors have occurred; execution is impossible. For example, if an invalid entry point has been specified, a return code of 0C is placed in register 15.
10	4	Terminal errors have occurred; the processing has terminated. For example, if the linkage editor cannot handle the blocking factor requested for SYSPRINT, a return code of 10 is placed in register 15.

Figure 16. Linkage Editor Return Codes

The programmer may use a return code to determine whether or not the load module is to be executed by using the condition parameter (COND) on the EXEC statement for the load module. The control program compares the return code with the values specified in the COND parameter, and the results of the comparisons are used to determine subsequent action. The COND parameter may be specified either in the JOB statement or the EXEC statement (see the publication JCL).

DD STATEMENTS

Every data set used by the linkage editor must be described with a DD statement. Each DD statement must have a name, unless data sets are concatenated. The DD statements for data sets required by the linkage editor have preassigned names; those for additional input data sets have user-assigned names; those for concatenated data sets (after the first) have no names.

In addition to the name, the DD statement provides the control program with information about the input/output device on which the data set resides, and a description of the data set itself. All of the job control language facilities for device description are available to the users of the linkage editor.

Besides information about the device, the DD statement also contains a data set description which includes the data set name and its disposition. Information for the data control block (DCB) may also be given.

General information pertinent to the linkage editor on the data set name and DCB information follows; information on disposition is given in the discussion for each data set.

DATA SET NAME: The linkage editor uses either sequential or partitioned data sets. For sequential data sets, only the name of the data set is specified; for partitioned data sets, the member name must also be specified either on the DD statement or with a control statement.

When input data sets are passed from a previous job step, or when the output load module is being tested, a recommended practice is to use temporary data set names (that is, &&dsname). Use of temporary names ensures that there are no duplicate data sets with out-of-date modules. A data set with a temporary name is automatically deleted at the end of the job. When a module is to be stored permanently, a data set name without ampersands is used.

DCB INFORMATION: Before a data set can be used for input, information describing the data set must be placed in the data control block (DCB). If this information does not exist in the DCB or header label, or if no labels are used (magnetic tape does not require labels), the programmer must specify it in the DCB parameter on the DD statement.

Record format (RECFM), logical record size (LRECL), and block size (BLKSIZE) subparameters of the DCB parameter are discussed as they apply to the linkage editor. Specific information on each as it applies to the linkage editor data sets is given in the description of the data set later in this section. Other DCB information (tape recording technique, density, and so forth) is described in the publication JCL.

Record Format: The following record formats are used with the linkage editor:

- F The records are fixed length.
- FB The records are fixed length and blocked.
- FBA The records are fixed length, blocked, and contain American National Standards Institute (ANSI) control characters.
- FBS The records are fixed length, blocked, and written in standard blocks.
- FA The records are fixed length and contain ANSI control characters.
- FS The records are fixed length and written in standard blocks.

- U The records are undefined length.
- UA The records are undefined length and contain ANSI control characters.

A record format of FS or FBS must be used with caution. All blocks in the data set must be the same size. This size must be equal to the specified block size. A truncated block can occur only as the last block in the data set.

Note: Track overflow is never used by the linkage editor. When moving or copying load modules, it is recommended that the track overflow feature not be used on the target data set, as errors may occur in fetching the load modules for execution.

LOGICAL RECORD AND BLOCK SIZE: Blocking is allowed for input object module data sets and the diagnostic output data set. The blocking factors used to determine buffer allocations are 5, 10, and 40. The BLKSIZE must therefore be a multiple of LRECL. See the description of blocking factors in the discussion of the SIZE option.

When the DCBS option is specified, a block size should be specified for the output load module library (see "SYSMOD DD Statement" on page 58).

LINKAGE EDITOR DD STATEMENTS

The linkage editor uses six data sets; of these, four are required. The DD statements for these data sets must use the preassigned ddnames given in Figure 17. The descriptions that follow give pertinent device and data set information for each linkage editor data set.

Data Set	ddname	Required
Primary input data set	SYSLIN	Yes
Automatic call library	SYSLIB	Only if the automatic library call mechanism is used
Intermediate data set	SYSUT1	Yes
Diagnostic output data set	SYSPRINT	Yes
Output module library	SYSMOD	Yes
Alternate output data set	SYSTEM	Only if the TERM option is specified

Figure 17. Linkage Editor ddnames

SYSLIN DD Statement

The SYSLIN DD statement is always required; it describes the primary input data set that can be assigned to a direct access device, a magnetic tape unit, or the card reader. The data set may be either sequential or partitioned; in the latter case, a member name must be specified.

If SYSLIN is assigned to a card reader or "pseudo card reader," input records must be unblocked and 80 bytes long. (A pseudo card reader is defined as input from a tape or a direct access device in card reader mode.)

This data set must contain object modules and/or control statements. Load modules used in the primary input data set are considered a severity 4 error.

The recommended disposition for the primary input data set is SHR or OLD.

The DCB requirements are shown in Figure 18.

LRECL	BLKSIZE	RECFM
80	80	F,FS
80	400,800,3200 ¹	FB,FBS

¹ These are the maximum block sizes allowed for each of the optimal blocking factors (5, 10, and 40). Which maximum is applicable depends on the value given to value1 and value2 of the SIZE option.

Figure 18. DCB Requirements for Object Module and Control Statement Input

SYSLIB DD Statement

The SYSLIB DD statement is required when the automatic library-call mechanism is to be used. This DD statement describes the automatic call library, which must be assigned to a direct access device. The data set must be partitioned, but member names should not be specified.

The recommended disposition for the call library is SHR or OLD.

If concatenated call libraries are used, object and load module libraries must not be mixed. If only object modules are used, the call library may also contain control statements.

The DCB requirements for object module call libraries are given in Figure 18. The DCB requirement for load module call libraries is a record format of U; the block size used for storage allocation is equal to the maximum for the device used, not the record read. Note that the linkage editor recognizes object and load module call libraries solely from their record format, and not from the data within them.

This data set must not be assigned to SYSOUT.

SYSUT1 DD Statement

The SYSUT1 DD statement is always required; it describes the intermediate data set, which is a sequential data set assigned to a direct access device. Space must be allocated for this data set, but the DCB requirements are supplied by the linkage editor.

SYSPRINT DD Statement

The SYSPRINT DD statement is always required; it describes the diagnostic output data set, which is a sequential data set assigned to a printer or to an intermediate storage device. If an intermediate storage device is used, the data records contain a carriage control character as the first byte.

The usual specification for this data set is SYSOUT=A. The programmer may assign a block size. The record format assigned

by the linkage editor depends on whether blocking is used or not.

Figure 19 shows the DCB requirements for SYSPRINT. The only information that can be supplied by the programmer is the block size.

LRECL	BLKSIZE	RECFM
121	121	FA
121	$n \times 121$ where n is less than or equal to 40	FBA

Note: The value specified for BLKSIZE, either on the DCB parameter of the SYSPRINT DD statement or in the DSCB (data set control block) of an existing data set, must be a multiple of 121; if it is not, the linkage editor issues a message to the operator's console and terminates processing.

Figure 19. DCB Requirements for SYSPRINT

SYSLMOD DD Statement

The SYSLMOD DD statement is always required; it describes the output module library, which must be a partitioned data set assigned to a direct access device.

A member name may be specified on the SYSLMOD DD statement. If a member name is specified, it is used only if a name was not specified on a NAME control statement. This member name must conform to the rules for the name on the NAME control statement. This would imply the replacement of an identically named member in the output load module library, if one exists.

If SYSLMOD is to be referenced by an INCLUDE statement, the member name on the DD statement, if present, must be the name of an existing member.

If the member is to replace an identically named member in an existing library, the disposition should be OLD or SHR. If the member is to be added to an existing library, the disposition should be MOD, OLD, or SHR. If no library exists and the member is the first to be added to a new library, the disposition should be NEW or MOD. If the member is to be added to an existing library that may be used concurrently in another region or partition, the disposition should be SHR.

The record format U is assigned by the linkage editor. See "Appendix D. Loader Storage Considerations" on page 187.

Procedures used by the linkage editor to assign block size are:

1. If the data set is new:
 - a. Without the DCBS option specified:
 - The DSCB (data set control block) reflects the maximum block size available for the device type if it is not restricted by value2 of the size parameter.
 - If SCTR is specified, the block size is 1024.

- b. With the DCBS option specified, the DSCB block size is the smaller of:
 - The maximum track size for the device.
 - The value of the BLKSIZE subparameter on the DCB parameter of the SYSLMOD DD statement.
 - The actual output buffer length (half the number specified for value2 if the size option was utilized).
- c. The minimum DSCB block size is 256 without the SCTR option specified and 1024 with the SCTR option.
2. For preallocated data sets not previously opened, a block size is assigned as for new data sets.
3. When the DSCB block size already exists (not a new or preallocated data set) and the SCTR option is specified, 1024 is used.
4. When the DSCB block size already exists and the DCBS or SCTR option is not specified, the larger of the existing block sizes or 256 is used.
5. See "DCBS Option" on page 50 for the procedure when the DSCB block size exists and the DCBS option is specified.

Note: When a new data set is created at linkage editor time without the DCBS option specified, the DSCB reflects the maximum block size available for the device type.

If the SYSLMOD DD statement is used as a source of load module input, the SYSLMOD data set is read with a record format of U in all cases.

In the following example, the SYSLMOD DD statement specifies a permanent library on an IBM 3380 Disk Storage Device:

```
//SYSLMOD DD DSNAME=USERLIB(TAXES),DISP=MOD,
// UNIT=3380,...
```

The linkage editor assigns a record format of U, and a maximum logical record and block size of 32760 bytes, the maximum for the sequential access method. However, consider the following example:

```
//LKED EXEC PGM=HEWL,PARM='XREF,DCBS'
.
//SYSLMOD DD DSNAME=USERLIB(TAXES),DISP=MOD,
// UNIT=3380,DCB=(BLKSIZE=13030),...
```

The linkage editor still assigns a record format of U, but the logical record and block size are now 13030 bytes rather than 32760 bytes, because of the use of the DCBS option.

SYSTEM DD Statement

The SYSTEM DD statement is optional; it describes a data set that is used only for numbered error/warning messages. Although intended to define the terminal data set when the linkage editor is being used under the Time Sharing Option (TSO) of MVS, the SYSTEM DD statement can be used in any environment to define a data set consisting of numbered error/warning messages that supplements the SYSPRINT data set.

SYSTEM output is defined by including a SYSTEM DD statement and specifying TERM in the PARM field of the EXEC statement. When SYSTEM output is defined, numbered messages are then written to both the SYSTEM and SYSPRINT data sets.

The following example shows how the SYSTEM DD statement could be used to specify the system output unit:

```
//SYSTEM DD SYSOUT=A
```

The DCB requirements for SYSTEM (LRECL=121, BLKSIZE=121, and RECFM=FBA) are supplied by the linkage editor. If necessary, the linkage editor will modify the DSCB (data set control block) of an existing data set to reflect these values.

ADDITIONAL DD STATEMENTS

Each ddname specified on an INCLUDE or a LIBRARY control statement must also be described with a DD statement. These DD statements describe sequential or partitioned data sets, assigned to magnetic tape units or direct access devices (not pseudo card readers).

The ddnames are specified by the user with any other necessary information. The DCB requirements for these data sets are shown in Figure 20.

	LRECL	BLKSIZE	RECFM
Include Control Statement			
Object modules and/or control statements	80	80	F,FS
Load modules	Maximum for device, or one-half of <u>value2</u> , whichever is smaller	Equal to LRECL	U
Library Control Statement			
Object modules and/or control statements	80 80	80 400,800,3200 ¹	F,FS FB,FBS
Load Modules	Maximum for device, or one-half of <u>value2</u> , whichever is smaller	Equal to LRECL	U

Figure 20. DCB Requirements for Additional Input Data Sets

Note to Figure 20:

¹ These are the maximum block sizes allowed for each of the optimal blocking factors (5, 10, 40). Which maximum is applicable depends on the values given to value1 and value2 of the SIZE option.

When concatenated data sets are included, each data set must contain records of the same format, record size, and block size. If the data sets reside on magnetic tape, the tape recording technique and density must also be identical.

If the SYSLMOD DD statement is used as a source of load module input, the SYSLMOD data set is read with a record format of U in all cases.

SIZE PARAMETER GUIDELINES

This section gives guidelines for determining appropriate SIZE parameter values for a linkage editor job step.

First—determine Value2 of the SIZE parameter.

$$\text{Value2} = [6K | 6144 | f | g | (a+b) | (c \times d) | (c \times e)]$$

where:

- a is the length of the load module to be built.
- b is 0, if the length of the load module to be built is < 40K bytes.
is 4K, if the length of the load module to be built is ≥ 40K bytes.
- c is an integer equal to or greater than 2, such that $c \times d$ or $c \times e$ is ≤ 999999 or 9999K bytes (c is the integer that represents the number of buffers to be reserved for SYSLMOD).
- d is the track capacity of the SYSLMOD device, or 32760, whichever is larger.
- e is the block size of the SYSLMOD data set.
- f is the length of the largest text record in load module input.
- g is the track capacity of the SYSUT1 device, or 32760, whichever is larger.

Selecting the largest of the above parameters provides optimal results.

Second—determine Value1 of the SIZE parameter.

$$\text{Value1} = h + j + k$$

Value1 must range between h and 9999K or 9999999

where:

- h = 96K
- j is the excess of Value2 over 6K
- k is the additional storage required to support the blocking factor for SYSLIN, object module libraries, and SYSPRINT:

Blocking Factor	K (Bytes)
5 to 1	0
10 to 1	18
40 to 1	28

Third—determine the REGION parameter.

REGION = Equal to or greater than Value1

CATALOGED PROCEDURES

To facilitate the operation of the system, the control program allows the programmer to store EXEC and DD statements under a unique member name in a procedure library. Such a series of job control language statements is called a cataloged procedure. These job control language statements can be recalled at any time to specify the requirements for a job. To request this procedure, the programmer places an EXEC statement in the input stream. This EXEC statement specifies the unique member name of the procedure desired.

The specifications in a cataloged procedure can be temporarily overridden, and DD statements can be added. The information altered by the programmer is in effect only for the duration of the job step; the cataloged procedures themselves are not altered permanently. Any additional DD statements supplied by the programmer must follow those that override the cataloged procedure.

LINKAGE EDITOR CATALOGED PROCEDURES

Two linkage editor cataloged procedures are provided: a single-step procedure that link-edits the input and produces a load module (procedure LKED), and a two-step procedure that link-edits the input, produces a load module, and executes that module (procedure LKEDG). Many of the cataloged procedures provided for language translators also contain linkage editor steps. The EXEC and DD statement specifications in these steps are similar to the specifications in the cataloged procedures described in the following paragraphs.

Procedure LKED

The cataloged procedure named LKED is a single-step procedure that link-edits the input, produces a load module, and passes the load module to another step in the same job. The statements in this procedure are shown in Figure 21; the following text describes these statements.

```
//LKED      EXEC  PGM=HEWL,PARM='XREF,LIST,LET,NCAL',REGION=512K
//SYSPRINT  DD    SYSOUT=A
//SYSLIN    DD    DDNAME=SYSIN
//SYSLMOD   DD    DSNNAME=&&GOSET(GO),SPACE=(1024,(50,20,1)),
//          UNIT=SYSDA,DISP=(MOD,PASS)
//SYSUT1    DD    UNIT=(SYSDA,SEP=(SYSLMOD,SYSLIN)),
//          SPACE=(1024,(200,20))
```

Figure 21. Statements in the LKED Cataloged Procedure

STATEMENT NUMBERS: The 8-digit numbers on the right side of each statement (not shown in Figure 21) are used to identify each statement and would be used, for example, when permanently modifying the cataloged procedure with the system utility program IEBUPDTE. For a description of this utility program, see Utilities.

EXEC STATEMENT: The PARM field specifies the XREF, LIST, LET, and NCAL options. If the automatic library-call mechanism is to be used, the NCAL option must be overridden, and a SYSLIB DD statement must be added. Overriding and adding DD statements is discussed later in this section.

SYSPRINT STATEMENT: The SYSPRINT DD statement specifies the SYSOUT class A, which is either a printer or an intermediate storage device. If an intermediate storage device is used, American National Standard Institute control characters accompany the data to be printed.

SYSLIN STATEMENT: The specification of DDNAME=SYSIN allows the programmer to specify any input data set as long as it fulfills the requirements for linkage editor input. The input data set must be defined with a DD statement with the ddname SYSIN. This data set may be either in the input stream or reside on a separate volume.

If the data set is in the input stream, the following SYSIN statement is used:

```
//LKED.SYSIN DD *
```

If this SYSIN statement is used, it may be anywhere in the job step DD statements as long as it follows all overriding DD statements. The object module decks and/or control statements should follow the SYSIN statement, with a delimiter statement (/*) at the end of the input.

If the data set resides on a separate volume, the following SYSIN statement is used:

```
//LKED.SYSIN DD (parameters describing the input data set)
```

If this SYSIN statement is used, it may be anywhere in the job step DD statements as long as it follows all overriding DD statements. Several data sets may be concatenated, as described in "Chapter 3. Defining Input to the Linkage Editor" on page 23.

SYSMOD STATEMENT: The SYSMOD DD statement specifies a temporary data set and a general space allocation. The disposition allows the next job step to execute the load module. If the load module is to reside permanently in a library, these general specifications must be overridden.

SYSUT1 STATEMENT: The SYSUT1 DD statement specifies that the intermediate data set is to reside on a direct access device, but not the same device as either the SYSMOD or the SYSLIN data sets. Again, a general space allocation is given.

SYSLIB STATEMENT: Note that there is no SYSLIB DD statement. If the automatic library-call mechanism is to be used with a cataloged procedure, a SYSLIB DD statement must be added; also, the NCAL option in the PARM field of the EXEC statement must be negated.

INVOKING THE LKED PROCEDURE: To invoke the LKED procedure, code the following EXEC statement:

```
//stepname EXEC LKED
```

where stepname is optional and is the name of the job step.

The following example shows a sample JCL sequence for using the LKED procedure in one step to link-edit object modules to produce a load module, then execute the load module in a subsequent step.

```

//LESTEP      EXEC    LKED
              (Overriding and additional DD statements for the LKED step)
//LKED.SYSIN DD      *
              (Object module decks and/or control statements)
//EXSTEP      EXEC    PGM=* .LESTEP.LKED.SYSLMOD
              (DD statements and data for load module execution)
/*           (If data is supplied for the execution step)

```

Note: LESTEP invokes the LKED procedure and EXSTEP executes the load module produced by LESTEP.

Procedure LKEDG

The cataloged procedure named LKEDG is a two-step procedure that link-edits the input, produces a load module, and executes that load module. The statements in this procedure are shown in Figure 22. The two steps are named LKED and GO. The specifications in the statements in the LKED step are identical to the specifications in the LKED procedure.

```

//LKED      EXEC    PGM=HEWL,PARM='XREF,LIST,NCAL',REGION=512K
//SYSPRINT  DD      SYSOUT=A
//SYSLIN    DD      DDNAME=SYSIN
//SYSLMOD   DD      DSN=*&&GOSSET(GO),SPACE=(1024,(50,20,1)),
//           UNIT=(SYSDA,DISP=(MOD,PASS))
//SYSUT1    DD      UNIT=(SYSDA,SEP=(SYSLMOD,SYSLIN)),
//           SPACE=(1024,(200,20))
//GO        EXEC    PGM=* .LKED.SYSLMOD,COND=(4,LT,LKED)

```

Figure 22. Statements in the LKEDG Cataloged Procedure

GO STEP: The EXEC statement specifies that the program to be executed is the load module produced in the LKED step of this job. This module was stored in the data set described on the SYSLMOD DD statement in that step. (If a NAME statement was used to specify a member name other than that used on the SYSLMOD statement, use the LKED procedure.)

The condition parameter specifies that the execution step is to be bypassed if the return code issued by the LKED step is greater than 4.

INVOKING THE LKEDG PROCEDURE: To invoke the LKEDG procedure, code the following EXEC statement:

```
//stepname EXEC LKEDG
```

where stepname is optional and is the name of the job step.

The following example shows a sample JCL sequence for using the LKEDG procedure to link-edit object modules, produce a load module, and execute that load module.

```

//TWOSTEP EXEC LKEDG.
    (Overriding and additional DD statements for the LKED step)
//LKED.SYSIN DD *
    (Object module decks and/or control statements)
/*
    (DD statements for the GO step)
//GO.SYSIN DD *
    (Data for the GO step)
/*

```

OVERRIDING CATALOGED PROCEDURES

The programmer may override any of the EXEC or DD statement specifications in a cataloged procedure. These new specifications remain in effect only for the duration of the job step. For a detailed description of overriding cataloged procedures, see the publication JCL.

Overriding the EXEC Statement

The EXEC statement in a cataloged procedure is overridden by specifying the changes and additions on the EXEC statement that invokes the cataloged procedure. The stepname should be specified when overriding the EXEC statement parameters.

For example, the REGION parameter can be increased as follows:

```
//LESTEP EXEC LKED,REGION.LKED=136K
```

The rest of the specifications on the EXEC statement of procedure LKED remain in effect.

If the PARM field is to be overridden, all the options specified in the cataloged procedure are negated. That is, if XREF, LIST, or NCAL is desired when overriding the PARM field, it must be respecified. In the following example, the OVLY option is added and the NCAL option is negated:

```
//LESTEP EXEC LKED,PARM.LKED='OVLY,XREF,LIST'
```

As a result, the XREF and LIST options are retained, but the NCAL option is negated; when NCAL is negated, a SYSLIB DD statement must be added.

If you use the LKEDG procedure and want to execute the load module just built, an efficient way is to specify the parameter LET in the LKED step and invoke the LKEDG procedure with the following EXEC statement:

```

//stepname EXEC LKEDG,PARM.LKED='XREF,LIST,NCAL,LET',
// COND.GO=(8,LT,LKED)

```


Overriding DD Statements

Each DD statement that is used to override a DD statement in the LKED step of either the LKED procedure or the LKEDG procedure must begin with `//LKED.ddname...`

Any of the DD statements in the cataloged procedures can be overridden as long as the overriding DD statements are in the same order as they appear in the procedure. If any DD statements are not overridden, or overriding DD statements are included but are not in sequence, the specifications in the cataloged procedure are used.

Only those parameters specified on the overriding DD statement are affected; the rest of the parameters remain as specified in the procedure. In the following example, the output load module is to be placed in a permanent library:

```
//LIBUPDTE EXEC LKED
//LKED.SYSLMOD DD DSN=LOADLIB(PAYROLL),DISP=OLD
//LKED.SYSIN DD DSN=OBJMOD,DISP=(OLD,DELETE)
```

Unit and volume information should be given if these data sets are not cataloged.

As a result of the statements in the example, the LKED procedure is used to process the object module in the OBJMOD data set. The output load module is stored in the data set LOADLIB with the name PAYROLL. The SPACE parameter on the SYSLMOD DD statement and the other specifications in the procedure remain in effect.

ADDING DD STATEMENTS

DD statements for additional data sets can be supplied when using cataloged procedures. These additional DD statements must follow any overriding DD statements.

Each additional DD statement for the LKED step must begin with `//LKED.ddname...`; for the GO step, it must begin with `//GO.ddname...`

In the following example, the automatic library-call mechanism is to be used along with the LKEDG procedure:

```
//CPSTEP EXEC LKEDG,PARM=LKED='XREF,LIST'
//LKED.SYSLMOD DD DSN=LOADLIB(TESTER),DISP=OLD,...
//LKED.SYSLIB DD DSN=SYLI.PL1LIB,DISP=SHR
//LKED.SYSIN DD *
```

(Object module decks and/or control statements).

```
/*
//GO.SYSIN DD *
```

(Data for execution step)

```
/*
```

The NCAL option is negated, and a SYSLIB DD statement is added between the overriding SYSLMOD DD statement and the SYSIN DD statement.

CHAPTER 5. SPECIFYING AN OPERATION WITH CONTROL STATEMENTS

This chapter summarizes the linkage editor control statements. The description of each statement includes:

- What the statement does
- The format of the statement
- Placement of the statement in the input
- Notes on use, if any
- One or more examples that include job control language statements, when necessary

The control statements are described in alphabetic order. Before using this chapter, the user should be familiar with the following information on general format, format conventions, and placement.

General Format

Each linkage editor control statement specifies an operation and one or more operands. Nothing must be written preceding the operation, which must begin in or after column 2. The operation must be separated from the operand by one or more blanks.

A control statement can be continued on as many cards as necessary by terminating the operand at a comma, and by placing a nonblank character in column 72 of the card. Continuation must begin in column 16 of the next card. A symbol cannot be split; that is, it cannot begin on one card and be continued on the next.

Format Conventions

The following conventions are used in the formats to describe the coding of the linkage editor control statements:

- **Boldface** type indicates the exact characters to be entered. Such items must be entered exactly as illustrated (in uppercase, if applicable).
- Lowercase underscored type specifies fields to be supplied by the user.
- Other punctuation (parentheses, commas, spaces, and so forth) must be entered as shown.
- Braces { } indicate a choice of entry; unless a default is indicated, you must choose one of the entries.
- Brackets [] indicate an optional field or parameter.
- An ellipsis (...) indicates that multiple entries of the type immediately preceding the ellipsis are allowed.
- Items separated by a vertical bar (|) represent alternative items. No more than one of the items may be selected.

Placement Information

Linkage editor control statements are placed before, between, or after modules. They can be grouped, but they cannot be placed within a module. However, specific placement restrictions may be imposed by the nature of the functions being requested by the control statement. Any placement restrictions are noted.

ALIAS Statement

The ALIAS statement specifies additional names for the output library member, and can also specify names of alternative entry points. Up to 16 names can be specified on one ALIAS statement, or separate ALIAS statements for one library member. The names are entered in the directory of the partitioned data set in addition to the member name.

FORMAT: The format of the ALIAS statement is:

ALIAS	{ <u>symbol</u> <u>external_name</u> }
-------	------------------------------------------

symbol

specifies an alternate name for the load module. When the module is executed, the main entry point is used as the starting point for execution.

external_name

specifies a name that is defined as a control section name or entry name in the output module. When the module is called for execution, execution begins at the external name referred to.

PLACEMENT: An ALIAS statement can be placed before, between, or after object modules or other control statements. It must precede a NAME statement used to specify the member name, if one is present.

Notes:

1. In an overlay program, an external name specified by the ALIAS statement must be in the root segment.
2. No more than 16 alias names can be assigned to one output module.
3. Each alias specified for a load module is retained in the directory entry for the module; the linkage editor does not delete an old alias. Therefore, each alias that is specified must be unique; assigning the same alias to more than one load module can cause incorrect module references.
4. Obsolete alias names should be deleted from the PDS directory using a system utility such as IEHPRGM, to avoid future name conflicts.
5. If the replace option is in effect for the output load module (that is, the load module built in this link-edit does or may replace an identically named load module in the output module library), the replace option is in effect for each ALIAS name for the load module as well as for the primary name.

EXAMPLE: An output module, ROUT1, is to be assigned two alternate entry points, CODE1 and CODE2. In addition, calling modules have been written using both ROUT1 and ROUTONE to refer to the output module. Rather than correct the calling modules, an alternative library member name is also assigned.

```
ALIAS      CODE1, CODE2, ROUTONE
NAME      ROUT1
```

Because CODE1 and CODE2 are entry names in the output module, execution begins at the point referred to when these names are used to call the module. The modules that call the output module with the name ROUTONE now correctly refer to ROUT1 at its main entry point. The names CODE1, CODE2, and ROUTONE appear in the library directory along with ROUT1.

CHANGE Statement

The CHANGE statement causes an external symbol to be replaced by the symbol in parentheses following the external symbol. The external symbol to be changed can be a control section name, an entry name, or an external reference. More than one such substitution may be specified in one CHANGE statement.

FORMAT: The format of the CHANGE statement is:

CHANGE	<u>externalsymbol</u> (<u>newsymbol</u>) [, <u>externalsymbol</u> (<u>newsymbol</u>)]...
--------	-------------------------------------------------------------------------------------------------

externalsymbol

is the control section name, entry name, or external reference that is to be changed.

newsymbol

is the name to which the external symbol is to be changed.

PLACEMENT: The CHANGE control statement must be placed immediately before either the module containing the external symbol to be changed, or the INCLUDE control statement specifying the module. The scope of the CHANGE statement is across the immediately following module (object module or load module); the END record in the immediately following object module or the end-of-module indication in the immediately following load module delimits the scope of the CHANGE statement.

Notes:

1. External references from other modules to a changed control section name or entry name remain unresolved unless further action is taken.
2. If the external symbol specified on the CHANGE statement is misspelled, the symbol will not be changed. Linkage editor output, such as the cross-reference listing or module map, can be used to verify each change.
3. When a REPLACE statement that deletes a control section is followed by a CHANGE statement with the same control section name, unpredictable results will occur.

EXAMPLE 1: Two control sections in different modules have the name TAXROUT. Because both modules are to be link-edited together, one of the control section names must be changed. The module to be changed is defined with a DD statement named OBJMOD. The control section name could be changed as follows:

```
//OBJMOD      DD      DSNAME=TAXES,DISP=(OLD,KEEP),...
//SYSLIN      DD      *
CHANGE TAXROUT(STATETAX)
INCLUDE OBJMOD
.
/*
```

As a result, the name of control section TAXROUT in module TAXES is changed to STATETAX.

EXAMPLE 2: A load module contains references to TAXROUT that must now be changed to STATETAX. This module is defined with a DD statement named LOADMOD. The external references could be changed at the same time the control section name is changed, as follows:

```
//OBJMOD      DD      DSNAME=TAXES,DISP=(OLD,DELETE),...
//LOADMOD     DD      DSNAME=LOADLIB,DISP=OLD,...
//SYSLIN      DD      *
               CHANGE TAXROUT(STATETAX)
               INCLUDE OBJMOD
               CHANGE TAXROUT(STATETAX)
               INCLUDE LOADMOD(INVENTORY)
               .
/*
```

As a result, control section name TAXROUT in module TAXES and external reference TAXROUT in module INVENTORY are both changed to STATETAX.

ENTRY Statement

The ENTRY statement specifies the symbolic name of the first instruction to be executed when the program is called by its module name for execution. An ENTRY statement should be used whenever a module is reprocessed by the linkage editor. If more than one ENTRY statement is encountered, the first statement specifies the main entry point; all other ENTRY statements are ignored.

FORMAT: The format of the ENTRY statement is:

ENTRY	<u>externalname</u>
-------	---------------------

externalname

is defined as either a control section name or an entry name in a linkage editor input module.

PLACEMENT: An ENTRY statement can be placed before, between, or after object modules or other control statements. It must precede the NAME statement for the module, if one is present.

Notes:

1. In an overlay program, the first instruction to be executed must be in the root segment.
2. The external name specified must be the name of an instruction, not a data name, if the module is to be executed.

EXAMPLE: In the following example, the main entry point is INIT1:

```
//LOADLIB DD DSNAME=LOADLIB,DISP=OLD,...
//SYSLIN DD *
ENTRY INIT1
INCLUDE LOADLIB(READ,WRITE)
.
.
ENTRY READIN
/*
```

INIT1 must be either a control section name or an entry name in the linkage editor input. The entry point specification of READIN is ignored.

EXPAND Statement

The EXPAND statement lengthens control sections or named common sections by a specified number of bytes.

FORMAT: The format of an EXPAND statement is

EXPAND	<code>name(XXXX) [,name(XXXX)]...</code>
--------	----------------------------------------------

name

is the symbolic name of a common section or control section whose length is to be increased.

XXXX

is the decimal number of bytes to be added to the length of a common section. The maximum is 4095 for each section indicated. Binary zeros will be added for an expanded control section.

The EXPAND statement is followed by a message, IEW0740, that indicates the number of bytes added to the control section and the offset, relative to the start of the control section, at which the expansion begins. The effective length of the expansion is given in hexadecimal and may be greater than the specified length if, after the specified expansion, padding bytes must be added for alignment of the next control section or named common section.

PLACEMENT: An EXPAND statement can be placed before, between, or after other control statements or object modules. However, the statement must follow the module containing the control or named common section to which it refers. If the control section or named common section is entered as the result of an INCLUDE statement, the EXPAND statement must immediately follow the INCLUDE statement.

Note: EXPAND should be used with caution so as not to increase the length of a program beyond its own design limitations. For example, if space is added to a control section beyond the range of its base register addressability, that space is unusable.

EXAMPLE: In the following example, EXPAND statements add a 250-byte patch area (initialized to zeros) at the end of control section CSECT1 and increase the length of named common section COM1 by 400 bytes.

```
//LKED      EXEC   PGM=HEWL
//SYSPRINT  DD     SYSOUT=A
//SYSUT1    DD     UNIT=SYSDA,SPACE=(TRK,(10,4))
//SYSLMOD   DD     DSNAME=PDSX,DISP=OLD
//SYSLIN    DD     DSNAME=&&LOADSET,DISP=(OLD,PASS),
//          DD     UNIT=SYSDA
//          DD     *
EXPAND      DD     CSECT1(250)
EXPAND      DD     COM1(400)
NAME       DD     MOD1(R)
/ *        DD     *
```


IDENTIFY Statement

The IDENTIFY statement specifies any data supplied by the user to be entered into the CSECT identification (IDR) records for a particular control section. The statement can be used either to supply descriptive data for a control section or to provide a means of associating system-supplied data with executable code.

FORMAT: The format of the IDENTIFY statement is:

IDENTIFY	<code><u>csectname('data')</u>[,<u>csectname('data')</u>]]...</code>
-----------------	----------------------------------------------------------------------

csectname

is the symbolic name of the control section to be identified.

data

specifies up to 40 EBCDIC characters of identifying information. The user may supply any information desired for identification purposes.

The rules of syntax for the operand field are:

1. No blanks or characters may appear between the left parenthesis and the leading single quotation mark nor between the trailing single quotation mark and the right parenthesis.
2. The data field consists of from 1 to 40 characters; therefore, a null entry must be represented, minimally, by a single blank.
3. Blanks may appear between the leading single quotation mark and the trailing single quotation mark. Each blank counts as 1 character toward the 40-character limit.
4. A single quotation mark between the leading quotation mark and the trailing quotation mark is represented by 2 consecutive quotation marks. The pair of quotation marks counts as 1 character toward the 40-character limit.
5. Any EBCDIC character may appear between the leading quotation mark and the trailing quotation mark. Each character counts as 1 character toward the 40-character limit.
6. The IDENTIFY statement may be continued; however, a whole operand must appear on a single card image and at least 1 whole operand must appear on each card image of the continued statement.
7. If a leading quotation mark is found, all characters are absorbed until a trailing quotation mark is found or the 40-character limit is exhausted.
8. Blanks may not appear between the CSECT name and the left parenthesis.
9. A blank following a left parenthesis terminates the operand field; a blank following a comma that terminates an operand also terminates the operand field of that card image.

PLACEMENT: An IDENTIFY statement can be placed before, between, or after other control statements or object modules. The IDENTIFY statement must follow the module containing the control section to be identified or the INCLUDE statement specifying the module.

Note: When two or more IDENTIFY statements specify the same CSECT name, only the last statement is effective.

EXAMPLE: In the following example, IDENTIFY statements are used to identify the source level of a control section, a PTF application to a control section, and the functions of several control sections.

```
//LKED      EXEC   PGM=HEWL
//SYSPRINT  DD     SYSOUT=A
//SYSUT1    DD     UNIT=SYSDA,SPACE=(TRK,(10,5))
//SYSLMOD   DD     DSNAME=LOADSET,DISP=OLD
//OLDMOD    DD     DSNAME=OLD.LOADSET,DISP=OLD
//PTFMOD    DD     DSNAME=PTF.OBJECT,DISP=OLD
//SYSLIN    DD     *

(input object deck for a control section named FORT)

  IDENTIFY  FORT('LEVEL 03')
  INCLUDE   PTFMOD(CSECT4)
  IDENTIFY  CSECT4('PTF99999')
  INCLUDE   OLDMOD(PROG1)
  IDENTIFY  CSECT1('I/O ROUTINE'),
            CSECT2('SORT ROUTINE'),
            CSECT3('SCAN ROUTINE')

/*
```

Execution of this example produces IDR records containing the following identification data:

- The name of the linkage editor that produced the load module, the linkage editor version and modification level, and the date of the current linkage editor processing of the module. This information is provided automatically.
- User-supplied data describing the functions of several control sections in the module, as indicated on the third IDENTIFY statement.
- If the language translator used supports IDR, the identification records produced by the linkage editor also contain the name of the translator that produced the object module, its version and modification level, and the data of compilation.

The IDR records created by the linkage editor can be referenced by using the LISTIDR function of the service aid program AMBLIST. For instructions on how to use AMBLIST, see Service Aids.

INCLUDE Statement

The INCLUDE statement specifies sequential data sets and/or libraries that are to be sources of additional input for the linkage editor. INCLUDE statements are processed in the order in which they appear in the input. However, the sequence of data sets and modules within the output load module does not necessarily follow the order of the INCLUDE statements. If the order of the CSECTs within the module is significant, the user must specify the desired sequence by using order cards.

FORMAT: The format of the INCLUDE statement is:

INCLUDE	<u>ddname</u> [(<u>membername</u> [,...])] [<u>ddname</u> [(<u>membername</u> [,...])]]...
----------------	-----------------------------------------------------------------------------------------------

ddname is the name of a DD statement that describes either a sequential or a partitioned data set to be used as additional input to the linkage editor. For a sequential data set, ddname is all that must be specified. For a partitioned data set, at least one member name must also be specified.

membername is the name of or an alias for a member of the library defined in the specified DD statement. The membername must not be specified again on the DD statement.

PLACEMENT: An INCLUDE statement can usually be placed before, between, or after object modules or other control statements. However, when link-editing the nucleus, any ORDER statements used should precede the INCLUDE statements.

Note: A NAME statement in any data set specified in an INCLUDE statement is invalid; the NAME statement is ignored. All other control statements are processed.

EXAMPLE 1: In the following example, an INCLUDE statement specifies two data sets to be the input to the linkage editor:

```
//OBJMOD      DD      DSNAME=&&OBJECT,DISP=(OLD,DELETE)
//LOADMOD     DD      DSNAME=LOADLIB,DISP=SHR,...
.
.
//SYSLIN      DD      *
      INCLUDE OBJMOD,LOADMOD(TESTMOD,READMOD)
/*
```

Note that a DD statement must be supplied for every ddname specified in an INCLUDE statement.

EXAMPLE 2: Two separate INCLUDE statements could have been used in the preceding example, as follows:

```
INCLUDE OBJMOD
INCLUDE LOADMOD(TESTMOD,READMOD)
```

INSERT Statement

The INSERT statement repositions a control section from its position in the input sequence to a segment in an overlay structure. However, the sequence of control sections within a segment is not necessarily the order of the INSERT statements.

If a symbol specified in the operand field of an INSERT statement is not present in the external symbol dictionary, it is entered as an external reference. If the reference has not been resolved at the end of primary input processing, the automatic library-call mechanism attempts to resolve it.

FORMAT: The format of the INSERT statement is:

INSERT	<u>csectname</u> ,...
--------	-----------------------

csectname

is the name of the control section to be repositioned. A particular control section can appear only once within a load module.

PLACEMENT: The INSERT statement must be placed in the input sequence following the OVERLAY statement that specifies the origin of the segment in which the control section is to be positioned. If the control section is to be positioned in the root segment, the INSERT statement must be placed before the first OVERLAY statement.

Note: Control sections that are positioned in a segment must contain all address constants to be used during execution unless:

- The A-type address constants are located in a segment in the path.
- The V-type address constants used to pass control to another segment are located in the path. If an exclusive reference is made, the V-type address constant must be in a common segment.
- The V-type address constants used with the SEGLD and SEGWT macro instructions are located in the segment.

EXAMPLE: The following INSERT (and OVERLAY) statements specify the overlay structure shown in Figure 23 on page 78:

```
//          EXEC   PGM=HEWL,PARM='OVLY,XREF,LIST'  
.  
.  
//SYSLIN     DD      *  
  INSERT CSA  
  INSERT CSB  
  OVERLAY ALPHA  
  INSERT CSC,CSD  
  OVERLAY ALPHA  
  INSERT CSE  
/*
```




Figure 23. Overlay Structure for INSERT Statement Example

LIBRARY Statement

The LIBRARY statement can be used to specify:

- Additional automatic call libraries, which contain modules used to resolve external references found in the program.
- Restricted no-call function: External references that are not to be resolved by the automatic library call mechanism during the current linkage editor job step.
- Never-call function: External references that are not to be resolved by the automatic library call mechanism during any linkage editor job step.

Combinations of these functions can be written in the same LIBRARY statement.

FORMAT: The format of the LIBRARY statement is:

LIBRARY	{ddname(membername[,...]) (externalreference[,...]) *(externalreference[,...])},...
---------	-------------------------------------------------------------------------------------------

ddname

is the name of a DD statement that defines a library.

membername

is the name of or an alias for a member of the specified library. Only those members specified are used to resolve references.

externalreference

is an external reference that may be unresolved after primary input processing. The external reference is not to be resolved by automatic library call.

*

indicates that the external reference is never to be resolved; if the * (asterisk) is missing, the reference is left unresolved only during the current linkage editor run.

PLACEMENT: A LIBRARY statement can be placed before, between, or after object modules or other control statements.

Notes:

1. If the unresolved external symbol is not a member name in the library specified, the external reference remains unresolved unless defined in another input module.
2. If the NCAL option is specified, the LIBRARY statement cannot be used to specify additional call libraries.
3. Members called by automatic library call are placed in the root segment of an overlay program, unless they are repositioned with an INSERT statement.
4. Specifying an external reference for restricted no-call or never-call by means of the LIBRARY statement prevents the external reference from being resolved by automatic inclusion of the necessary module from an automatic call library; it does not prevent the external reference from being resolved if the module necessary to resolve the reference is specifically included or is included as part of an input module.

EXAMPLE: The following example shows all three uses of the LIBRARY statement:

```
//          EXEC   PGM=HEWL,PARM='LET,XREF,LIST'  
//TESTLIB   DD     DSNAME=TEST,DISP=SHR,...  
           .  
           .  
//SYSLIN   DD     *  
           LIBRARY TESTLIB(DATE,TIME),(FICACOMP),*(STATETAX)  
/*
```

As a result, members DATE and TIME from the additional library TESTLIB are used to resolve external references. FICACOMP and STATETAX are not resolved; however, because the references remain unresolved, the LET option must be specified on the EXEC statement if the module is to be marked executable. In addition, STATETAX will not be resolved in any subsequent reprocessing by the linkage editor.

MODE Statement

The MODE statement specifies the residence mode for the output load module and/or the addressing mode for all the entry points into the load module (the main entry point, its true aliases, and all the alternate entry points).

FORMAT: The format of the MODE statement is as follows:

MODE	<u>modespec</u> (, <u>modespec</u>)
------	--------------------------------------

modespec

is either of the following:

- The designation of an addressing mode for the output load module by one of the following:
 - AMODE(24)
 - AMODE(31)
 - AMODE(ANY)
- The designation of residence mode for the output load module by one of the following:
 - RMODE(24)
 - RMODE(ANY)

PLACEMENT: The MODE control statement can be placed before, between, or after object modules or other control statements. It must precede the NAME statement for the module, if one is present.

Notes:

1. The residence mode assigned by the MODE control statement overrides the residence mode accumulated from the input control sections and private code. The residence mode assigned by the MODE control statement also overrides the residence mode assigned by the RMODE parameter in the PARM field of the EXEC statement.
2. The addressing mode assigned by the MODE control statement overrides the separate addressing modes found in the ESD data for the control sections within which the entry points are located. The addressing mode assigned by the MODE control statement overrides the addressing mode assigned by the AMODE parameter in the PARM field of the EXEC statement.
3. If more than one MODE control statement is encountered in the link-edit of a load module, the last valid mode specification is used. Likewise, if a mode specification occurs more than once within a MODE statement, the last valid mode specification is used.

4. If only one value, either AMODE or RMODE, is specified in the MODE control statement, the other value is implied according to the following table:

Value Specified	Value Implied
AMODE=24	RMODE=24
AMODE=31	RMODE=24
AMODE=ANY	RMODE=24
RMODE=24	see below
RMODE=ANY	AMODE=31

If only an RMODE of 24 is specified, no overriding AMODE value is assigned; instead, the AMODE value in the ESD data for the main entry point, a true alias, or an alternate entry point is used in generating its respective directory entry.

5. In generating a directory entry for either the main entry point, a true alias, or an alternate entry point, the linkage editor validates the combination of the AMODE value and the RMODE value, as specified by the user in the MODE control statement(s), according to the table below:

	RMODE=24	RMODE=ANY
AMODE=24	valid	invalid
AMODE=31	valid	valid
AMODE=ANY	valid	invalid

6. If the AMODE/RMODE combination resulting from the MODE control statement(s) is invalid, an error message is issued and the linkage editor ignores the MODE control statement(s) as the source of AMODE/RMODE data.

EXAMPLE: In the following example, an output load module, named NEWMOD, is created; it is given a true alias of TESTMOD; the residence mode for the load module is ANY; the addressing mode for both the main entry point, NEWMOD, and the true alias, TESTMOD, is 31.

```

//SYSLMOD DD DSN=TESTLOAD, DISP=MOD,...
//SYSLIN DD *
.
.
.
MODE AMODE(31),RMODE(ANY)
ALIAS TESTMOD
NAME NEWMOD
/*

```

NAME Statement

The NAME statement specifies the name of the load module created from the preceding input modules, and serves as a delimiter for input to the load module. As a delimiter, the NAME statement allows multiple load module processing in one linkage editor job step. The NAME statement can also indicate that the load module replaces an identically named module in the output module library.

FORMAT: The format of the NAME statement is:

NAME	<u>membername</u>[(R)]
-------------	-------------------------------

membername

is the name to be assigned to the load module that is created from the preceding input modules.

(R)

indicates that this load module replaces an identically named module in the output module library. If the module is not a replacement, the parenthesized value (R) should not be specified.

PLACEMENT: The NAME statement is placed after the last input module or control statement that is to be used for the output module.

Notes:

1. Any ALIAS statement used must precede the NAME statement.
2. A NAME statement found in a data set other than the primary input data set is invalid. The statement is ignored.

EXAMPLE: In the following example, two load modules, RDMOD and WRTMOD, are produced by the linkage editor in one job step:

```
//SYSLMOD DD DSNAME=AUXMODS,DISP=MOD,...
//NEWMOD DD DSNAME=&&WRTMOD,DISP=OLD
//SYSLIN DD DSNAME=&&RDMOD,DISP=OLD
// DD *
NAME RDMOD(R)
INCLUDE NEWMOD
NAME WRTMOD
/*
```

As a result, the first module is named RDMOD and replaces an identically named module in the output module library AUXMODS; the second module is named WRTMOD and is added to the library.

ORDER Statement

The ORDER statement indicates the sequence in which control sections or named common areas appear in the output load module. The control sections or named common areas appear in the sequence in which they are specified on the ORDER statement. When multiple ORDER statements are used, their sequence further determines the sequence of the control sections or named common areas in the output load module; those named on the first statement appear first, and so forth.

FORMAT: The format of the ORDER statement is:

ORDER	{ <u>common area name</u> [(P)] <u>csectname</u> [(P)]},...
-------	--------------------------------------------------------------

common area name

is the name of the common area to be sequenced.

csectname

is the name of the control section to be sequenced.

(P)

indicates that the starting address of the control section or named common area is to be on a page boundary within the load module. The control sections or common areas are aligned on 4K-byte page boundaries.

PLACEMENT: An ORDER statement can usually be placed before, between, or after object modules or other control statements. However, when link-editing the nucleus, any ORDER statements used should precede the INCLUDE statements.

Notes:

1. A control section or common area can be named on only one ORDER statement. If the same name is used more than once, except when it is the last operand on one ORDER statement and the first operand on the next, the name is ignored, as is the balance of the control statement on which it appears.
2. The control sections and common areas named as operands can appear in either the primary input or the automatic call library, or both.
3. If a control section or a named common area is changed by a CHANGE or REPLACE control statement and sequencing is desired, specify the new name on the ORDER statement. The ORDER statement refers to the control section by its new name.

EXAMPLE: In this example, the control sections in the load module LDMOD are arranged by the linkage editor according to the sequence specified on ORDER statements. The page boundary alignments and the control section sequence made as a result of these statements are shown in Figure 24 on page 85. Assume each control section is 1K byte in length.

JCL and Control Statements

```
.  
. .  
//SYSLMOD DD DSNAME=PVTLIB,DISP=OLD,...  
//SYSLIN DD *  
ORDER ROOTSEG(P),MAINSEG,SEG1,SEG2  
ORDER SEG3(P),ENTRY1  
ORDER FSTPART,SESECTA,SESECTB(P)  
INCLUDE SYSLMOD(LDMOD)  
/*
```

Output Load Module




Figure 24. Output Load Module for ORDER Statement Example. The control section name PART1 is changed by a CHANGE statement to FSTPART. The ORDER statement refers to the control section by its new name.

OVERLAY Statement

The OVERLAY statement indicates either the beginning of an overlay segment, or of an overlay region. Because a segment or a region is not named, the programmer identifies it by giving its origin (or load point) a symbolic name. This name is then used on an OVERLAY statement to signify the start of a new segment or region.

FORMAT: The format of the OVERLAY statement is:

OVERLAY	<u>symbol</u> (REGION)
---------	------------------------

symbol

is the symbolic name assigned to the origin of a segment. This symbol is not related to external symbols in a module.

(REGION)

specifies the origin of a new region.

PLACEMENT: The OVERLAY statement must precede the first module of the next segment, the INCLUDE statement specifying the first module of the segment, or the INSERT statement specifying the control sections to be positioned in the segment.

Notes:

1. The OVLY option must be specified on the EXEC statement when OVERLAY statements are to be used.
2. The sequence of OVERLAY statements should reflect the order of the segments in the overlay structure from top to bottom, left to right, and region by region.
3. No OVERLAY statement should precede the root segment.

EXAMPLE: The following OVERLAY and INSERT statements specify the overlay structure in Figure 25 on page 87.

```
//          EXEC   PGM=HEWL,PARM='OVLY,XREF,LIST'  
:  
//SYSLIN      DD    DSNAME=&&OBJ,...  
//           DD    *  
INSERT CSA  
OVERLAY ONE  
INSERT CSB  
OVERLAY TWO  
INSERT CSC  
OVERLAY TWO  
INSERT CSD  
OVERLAY ONE  
INSERT CSE,CSF  
OVERLAY THREE(REGION)  
INSERT CSH  
OVERLAY THREE  
INSERT CSI  
/*
```

REGION 1

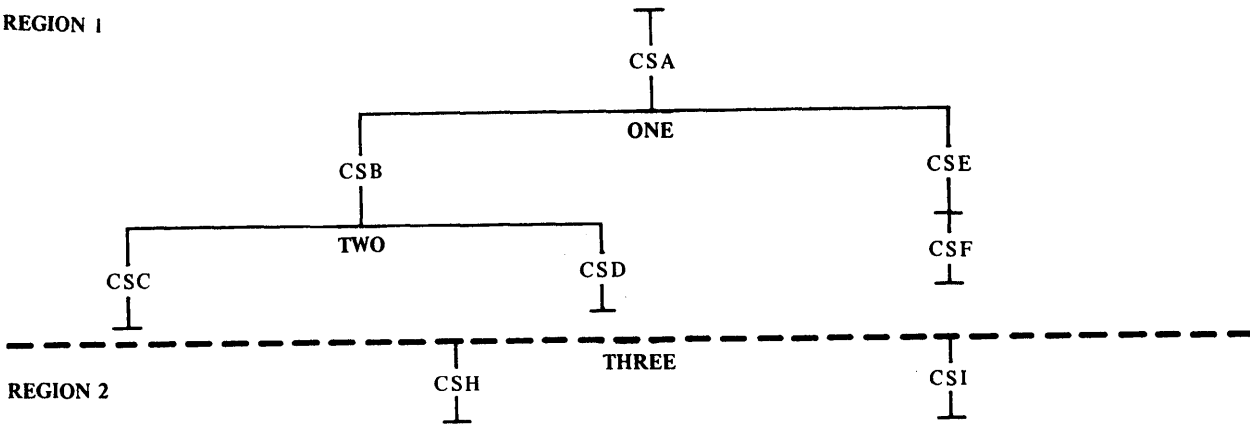


Figure 25. Overlay Structure for OVERLAY Statement Example

PAGE Statement

The PAGE statement aligns a control section or named common area on a 4K-byte page boundary in the load module.

FORMAT: The format of the PAGE statement is:

PAGE	{ <u>common area name</u> <u>csectname</u> },...
------	----------------------------------------------------

common area name

is the name of the common area to be aligned on a page boundary.

csectname

is the name of the control section to be aligned on a page boundary.

PLACEMENT: The PAGE statement can be placed before, between, or after object modules or other control statements.

Notes:

1. If a control section or a named common area is changed by a CHANGE or REPLACE control statement, and page alignment is wanted, specify the new name in the PAGE statement.
2. The control sections and common areas named as operands can appear in either the primary input or the automatic call library, or both.

EXAMPLE: In this example, the control sections in the load module LDMOD are aligned on page boundaries as specified in the following PAGE statement:

```
PAGE ALIGN,BNDRY4K,EIGHTK
```

The job control statements and linkage editor control statements as well as the output load module are shown in Figure 26 on page 89. Assume each control section is 3K bytes in length.

JCL And Control Statements

```
//LKED      EXEC  PGM=HEWL, PARM=, ...  
           .  
           .  
           .  
//SYSLMOD   DD    DSN=PVTLIB, DISP=OLD, ...  
//SYSLIN    DD    *  
           PAGE   ALIGN, BNDRY4K, EIGHTK  
           INCLUDE SYSLMOD(LDMOD)  
/*
```

Output Load Module




Figure 26. Output Load Module for PAGE Statement Example

REPLACE Statement

The REPLACE statement specifies one or more of the following:

- The replacement of one control section with another
- The deletion of a control section
- The deletion of an entry name

When a control section is replaced, all references within the input module to the old control section are changed to the new control section. Any external references to the old control section from other modules are unresolved unless changed.

When a control section is deleted, the control section name is also deleted from the external symbol dictionary, unless references are made to the control section from within the input module. If there are any such references, the control section name is changed to an external reference. External references from other modules to a deleted control section also remain unresolved.

When deleting an entry name, if there are any references to it within the same input module, the entry name is changed to an external reference.

FORMAT: The format of the REPLACE statement is:

REPLACE	{ <u>csectname-1</u> [(<u>csectname-2</u>)], <u>entryname</u> }
----------------	-------------------------------------------------------------------

csectname

is the name of a control section. If only csectname-1 is used, the control section is deleted; if csectname-2 is also used, the first control section is replaced with the second.

entryname

is the entry name to be deleted.

PLACEMENT: The REPLACE statement must immediately precede either (1) the module containing the control section or entry name to be replaced or deleted, or (2) the INCLUDE statement specifying the module. The scope of the REPLACE statement is across the immediately following module (object module or load module). The END record in the immediately following object module or the end-of-module indication in the load module terminates the action of the REPLACE statement. If the REPLACE statement is the last control statement in the SYSLIN data set, and there are unresolved external references to be resolved from SYSLIB, the REPLACE function operates on the first module from SYSLIB by an AUTO CALL.

Notes:

1. Unresolved external references are not deleted from the output module even though a deleted control section contains the only reference to a symbol.
2. When some but not all control sections of a separately assembled module are to be replaced, A-type address constants that refer to a deleted symbol will be incorrectly resolved, unless the entry name is at the same displacement from the origin in both the old and the new control sections.
3. If no INCLUDE statement follows the REPLACE statement, one module may be left out of AUTO CALL. Message IEW0132 is issued.
4. If the control section identified as csectname-1 (specified on the REPLACE statement) is misspelled, the control section

will not be replaced or deleted. Linkage editor output, such as the cross-reference listing and module map, can be used to verify each change.

EXAMPLE: In the following example, assume that control section INT7 is in member LOANCOMP and that control section INT8, which is to replace INT7, is in data set &&NEWINT. Also assume that control section PRIME in member LOANCOMP is to be deleted.

```
//NEWMOD      DD      DSNAME=&&NEWINT,DISP=(OLD,DELETE)
//OLDMOD      DD      DSNAME=PVTLIB,DISP=OLD,...
//SYSLIN      DD      *
  ENTRY MAINENT
  INCLUDE NEWMOD
  REPLACE INT7(INT8),PRIME
  INCLUDE OLDMOD(LOANCOMP)
/*
```

As a result, INT7 is removed from the input module described by the OLDMOD DD statement, and INT8 replaces INT7. All references to INT7 in the input module now refer to INT8. Any references to INT7 from other modules remain unresolved. If there are no references to PRIME in LOANCOMP, control section PRIME is deleted; the control section name is also deleted from the external symbol dictionary.

SETCODE Statement

The SETCODE statement assigns the specified authorization code to the output load module. The authorization code is placed in the directory entry for the output load module.

FORMAT: The format of the SETCODE statement is as follows:

SETCODE	AC(authorizationcode)
---------	-----------------------

authorizationcode

is 1 to 3 decimal digits specifying a value from 0 to 255.

PLACEMENT: A SETCODE statement can be placed before, between, or after object modules or other control statements. It must precede the NAME statement for the module, if one is present.

Notes:

1. The authorization code assigned by the SETCODE statement overrides the authorization code assigned by the AC parameter in the PARM field of the EXEC statement.
2. If more than one SETCODE statement is encountered in the link-edit of a load module, the last valid authorization code assigned is used.
3. The operand 'AC()' results in an authorization code of zero.

EXAMPLE: In the following example, an authorization code of 1 is assigned to the output load module MOD1.

```
//LKED      EXEC    PGM=HEWL
//SYSPRINT  DD      SYSOUT=A
//SYSUT1    DD      UNIT=SYSDA,SPACE=(TRK,(10,5))
//SYSLMOD   DD      DSNAME=SYS1.LINKLIB,DISP=OLD
//SYSLIN    DD      DSNAME=##LOADSET,DISP=(OLD,PASS)
//          DD      UNIT=SYSDA
//          DD      *
//          SETCODE AC(1)
//          NAME    MOD1(R)
/*
```

SETSSI Statement

The SETSSI statement specifies hexadecimal information to be placed in the system status index of the directory entry for the output module.

FORMAT: The format of the SETSSI statement is:

SETSSI	<u>XXXXXXXX</u>
--------	-----------------

XXXXXXXX

represents 8 hexadecimal characters (0 through 9 and A through F) to be placed in the 4-byte system status index of the output module library directory entry.

PLACEMENT: The SETSSI statement can be placed before, between, or after object modules or other control statements. If one is present, it must precede the NAME statement for the module.

Note: A SETSSI statement must be provided whenever an IBM-supplied load module is reprocessed by the linkage editor. If the statement is omitted, no system status index information is present.

CHAPTER 6. EDITING A CONTROL SECTION

The linkage editor performs editing functions either automatically or as directed by control statements. These editing functions provide for program modification on a control section basis. That is, they make it possible to modify a control section within an object or load module, without recompiling the entire source program.

The editing functions can modify either an entire control section or external symbols within a control section. Control sections can be deleted, replaced, or arranged in sequence; external symbols can be deleted or changed. (External symbols are control section names, entry names, external references, named common areas, or pseudoregisters.)

Whatever function is used, it is requested in reference to an input module. The resulting output load module reflects the request. That is, no actual change, deletion, or replacement is made to an input module. The requested alterations are used to control linkage editor processing (Figure 27).

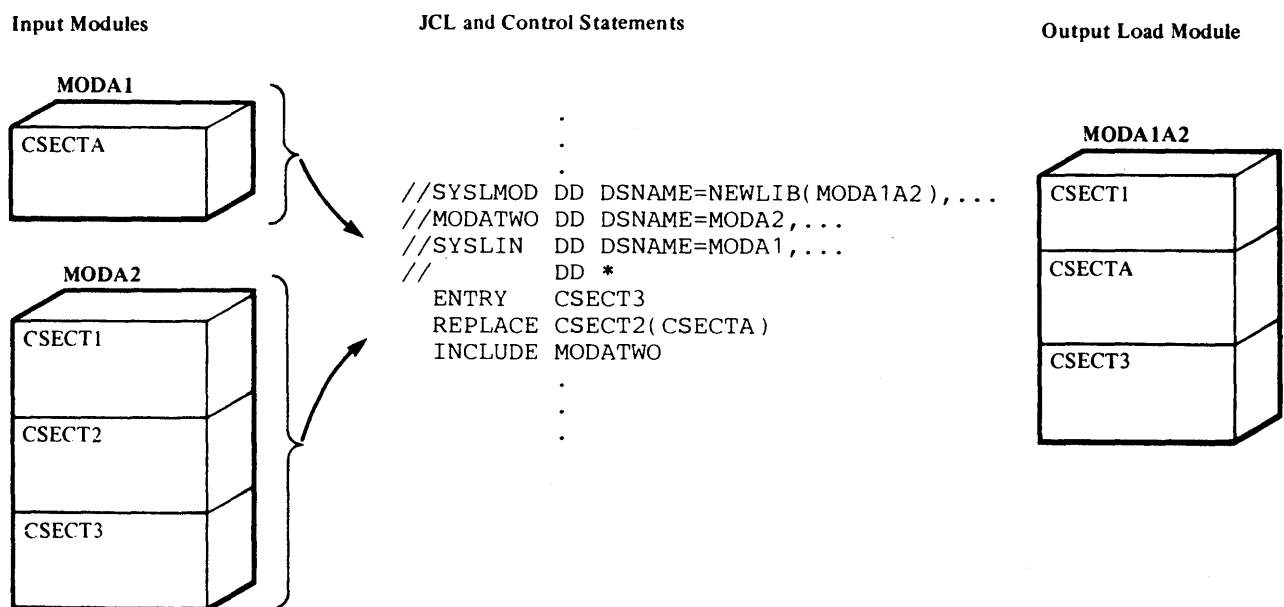


Figure 27. Editing a Module

Editing Conventions

In requesting editing functions, certain conventions should be followed to ensure that the specified modification is processed correctly. These conventions concern the following items:

- Entry points for the new module
- Placement of control statements
- Identical old and new symbols

ENTRY POINTS: Each time the linkage editor reprocesses a load module, the entry point for the output module should be specified in one of two ways:

- Through an ENTRY control statement.
- Through the assembler-produced END statement of an input object module, if one is present. If the entry point specified in the assembler-produced END statement is not defined in the object module, the entry name must be defined as an external reference.

The entry point assigned must be defined as an external name within the resulting load module.

PLACEMENT OF CONTROL STATEMENTS: The control statement (such as CHANGE or REPLACE) used to specify an editing function must precede either the module to be modified, or the INCLUDE statement that specifies the module. If an INCLUDE statement specifies several modules, the CHANGE or REPLACE statement applies only to the first module included.

IDENTICAL OLD AND NEW SYMBOLS: The same symbol should not appear as both an old external symbol and a new external symbol in one linkage editor run. If a control section is to be replaced by another control section with the same name, the linkage editor handles this automatically (see "Automatic Replacement" on page 98).

CHANGING EXTERNAL SYMBOLS

The linkage editor can be directed to change an external symbol to a new symbol while processing an input module. External references and address constants within the module automatically refer to the new symbol. External references from other modules to a changed external symbol must be changed with separate control statements.

Both the old and the new symbols are specified on either a CHANGE control statement or a REPLACE control statement. The use of the old symbol within the module determines whether the new symbol becomes a control section name, an entry name, or an external reference. The old symbol appears first, followed by the new symbol in parentheses.

The CHANGE control statement changes a control section name, an entry name, or an external reference. The REPLACE statement changes or deletes an entry name; if the symbols on a REPLACE statement are control section names, the entire control section is replaced or deleted (see "Replacing Control Sections" on page 97).

The CHANGE statement must immediately precede either the input module that contains the external symbol to be changed, or the INCLUDE statement that specifies the input module. The scope of the CHANGE statement is across the immediately following module (object module or load module). The END record in the immediately following object module or the end-of-module indication in the load module terminates the action of the CHANGE statement.

In the following example, assume that SUBONE is defined as an external reference in the input load module. A CHANGE statement is used to change the external reference to NEWMOD (Figure 28 on page 97).

```
//SYSLMOD      DD      DSNAME=PVTLIB,DISP=OLD,UNIT=3380,
//            VOLUME=SER=PVT002
//SYSLIN       DD      *
  ENTRY        BEGIN
  CHANGE       SUBONE(NEWMOD)
  INCLUDE      SYSLMOD(MAINROUT)
  NAME        MAINROUT(R)
/*
```

Input Module

JCL and Control Statements

Output Load Module

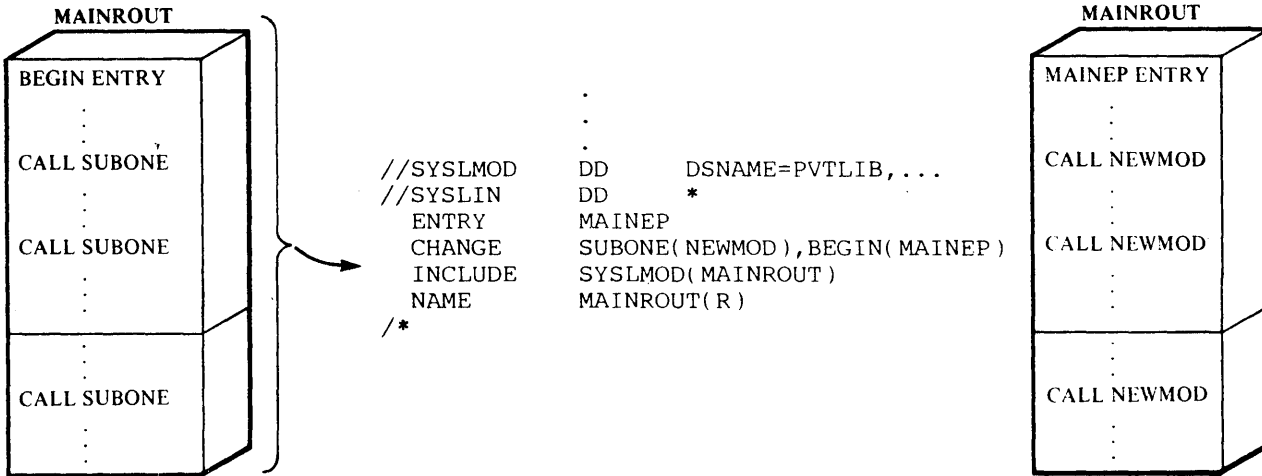


Figure 28. Changing an External Reference and an Entry Point

In the load module MAINROUT, every reference to SUBONE is changed to NEWMOD. Note also that the INCLUDE statement specifies a ddname of SYSLMOD. This allows a library to be used both as input and as the output module library.

More than one change can be specified on the same control statement. If, in the same example, the entry point is also to be changed, the two changes can be specified at once (see Figure 28).

```

//SYSLMOD DD DSNAME=PVTLIB,DISP=OLD,UNIT=3380,
// VOLUME=SER=PVT002
//SYSLIN DD *
ENTRY MAINEP
CHANGE SUBONE(NEWMOD),BEGIN(MAINEP)
INCLUDE SYSLMOD(MAINROUT)
NAME MAINROUT(R)
/*

```

The main entry point is now MAINEP instead of BEGIN. The ENTRY control statement specifies the new entry point, because this is the source of the name that is entered in the library directory entry for the load module's entry point.

REPLACING CONTROL SECTIONS

An entire control section can be replaced with a new control section. Control sections can be replaced either automatically or with a REPLACE control statement. Automatic replacement acts upon all input modules; the REPLACE statement acts only upon the module that follows it.

Notes:

1. Any CSECT identification (IDR) records associated with a particular control section are also replaced.
2. **For Assembler language programmers only:** When some but not all control sections of a separately assembled module are to be replaced, A-type address constants that refer to a deleted symbol will be incorrectly resolved unless the entry name is at the same displacement from the origin in both the old and the new control section. If all control sections of a separately assembled module are replaced, no restrictions apply.

AUTOMATIC REPLACEMENT

Control sections are automatically replaced if both the old and the new control section have the same name. The first of the identically named control sections processed by the linkage editor is made a part of the output module. All subsequent identically named control sections are ignored; external references to identically named control sections are resolved with respect to the first one processed. Therefore, to cause automatic replacement, the new control section must have the same name as the control section to be replaced, and must be processed before the old control section.

Caution: Automatic replacement applies to duplicate control section names only; if duplicate entry points exist in control sections with different names, a REPLACE control statement must be used to specify the entry point name. If a control section being automatically replaced contains unresolved external references and the control section replacing it does not, the parameter NCAL must be specified or the unresolved external references must be explicitly deleted using the REPLACE statement or marked for restricted no-call or never-call using the LIBRARY statement; otherwise, the unresolved external reference is retained.

NOTE ON OVERLAY PROGRAMS: When identically named control sections appear in modules being placed in an overlay structure, the second and any subsequent control sections with that name are ignored. This occurs whether the modules are in segments in the same path or in exclusive segments. Resolution of external references may therefore cause invalid exclusive references. Invalid exclusive references cause the linkage editor to mark the output module not executable unless the exclusive call (XCAL) option is specified on the EXEC statement (see "Chapter 4. Specifying JCL to Run a Linkage Editor Job" on page 37).

Example 1

An object module deck contains two control sections, READ and WRITE; member INOUT of library PVTLIB also contains a control section WRITE.

```
//SYSLMOD DD DSNAME=PVTLIB,DISP=OLD,UNIT=3380,
// VOLUME=SER=PVT002
//SYSLIN DD *

Object Deck for READ
Object Deck for WRITE

ENTRY READIN
INCLUDE SYSLMOD(INOUT)
NAME INOUT(R)
/*
```

The output load module contains the new READ control section, the new WRITE control section (replacing the old WRITE control section in member INOUT), and all remaining control sections from INOUT.

Example 2

A large load module named PAYROLL, originally written in COBOL, contains many control sections. Two control sections, FICA and STATETAX, were recompiled and passed to the linkage editor job step in the &&OBJECT data set. Then, by including the load module PAYROLL (a member of the partitioned data set LIB001) as well as the output of the language translator, the modified control sections automatically replace the identically named control sections (Figure 29 on page 100).

```
//SYSLMOD DD DSNAME=LIB002(PAYROLL),DISP=OLD,  
//          UNIT=3380,VOLUME=SER=LIB002  
//SYSLIB  DD DSNAME=SYS1.COBLIB,DISP=SHR  
//OLDLOAD DD DSNAME=LIB001,DISP=(OLD,DELETE),  
//          UNIT=3380,VOLUME=SER=LIB001  
//SYSLIN  DD DSNAME=&&OBJECT,DISP=(OLD,DELETE)  
//          DD *  
          INCLUDE OLDLOAD(PAYROLL)  
          ENTRY  INIT1  
/*
```

Input Modules

JCL and Control Statements

Output Load Module

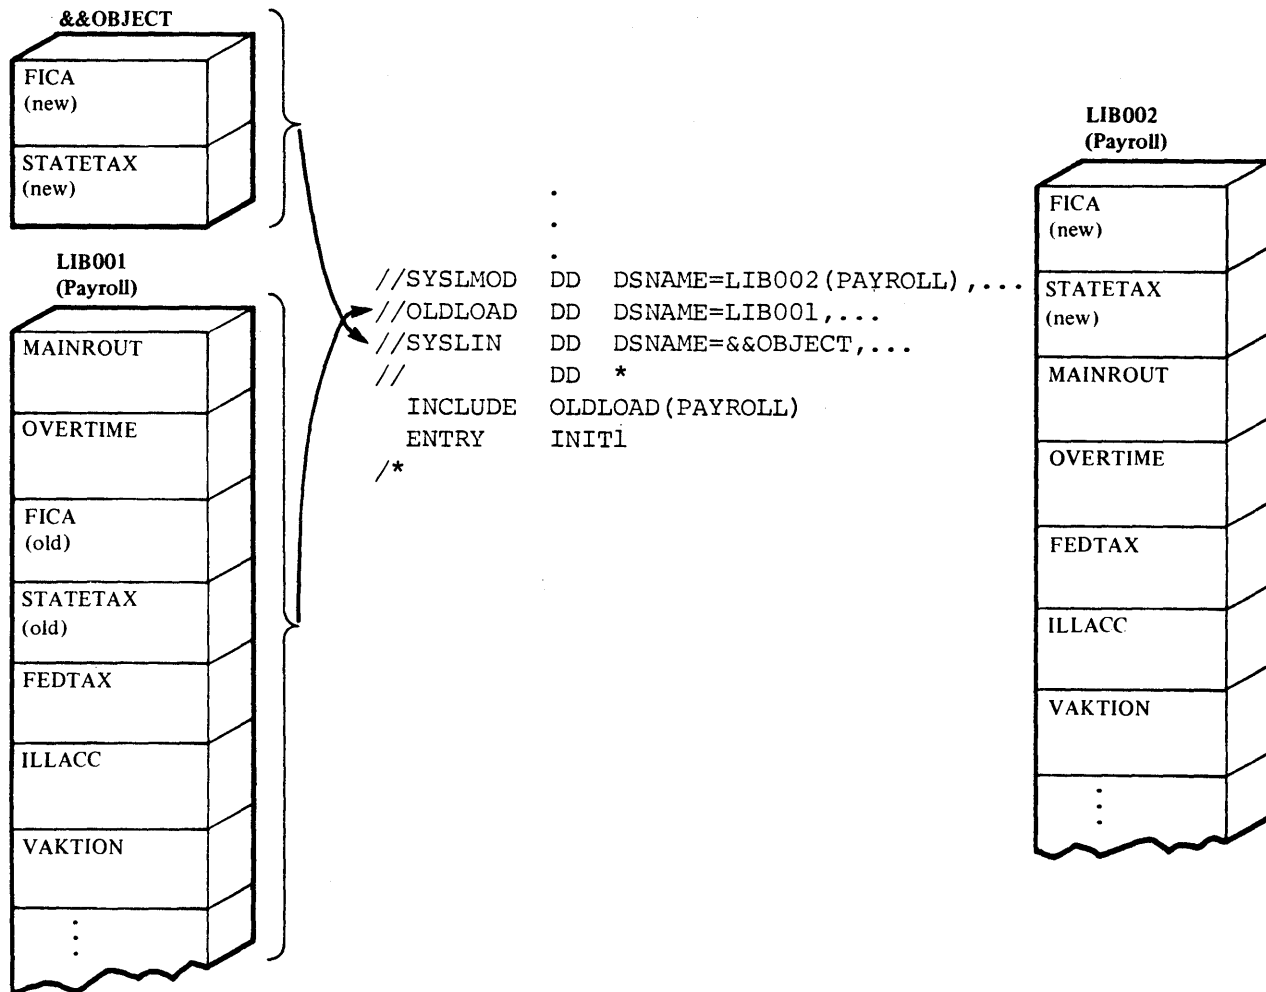


Figure 29. Automatic Replacement of Control Sections

The output module contains the modified FICA and STATETAX control sections and the rest of the control sections from the old PAYROLL module. The main entry point is INIT1, and the output module is placed in a library named LIB002. The COBOL automatic call library is used to resolve any external references that may be unresolved after the SYSLIN data sets are processed.

REPLACE STATEMENT

The REPLACE statement is used to replace control sections when the old and the new control sections have different names. The name of the old control section appears first, followed by the name of the new control section in parentheses. The REPLACE statement must precede either the input module that contains the control section to be replaced, or the INCLUDE statement that specifies the input module. The scope of the REPLACE statement is across the immediately following module (object module or load module). The END record in the immediately following object module or the end-of-module indication in the load module terminates the action of the REPLACE statement.

An external reference to the old control section from within the same input module is resolved to the new control section. An external reference to the old control section from any other module becomes an unresolved external reference unless one of the following occurs:

- The external reference to the old control section is changed to the new control section with a separate CHANGE control statement.
- The same entry name appears in the new control section or in some other control section in the linkage editor input.

In the following example, the REPLACE statement is used to replace one control section with another of a different name. Assume that the old control section SEARCH is in library member TBLESRCH, and that the new control section BINSRCH is in the data set &&OBJECT, which was passed from a previous step (Figure 30 on page 102).

```
//SYSLMOD DD DSNAME=SRCHRTN,DISP=OLD,UNIT=3380,  
// VOLUME=SER=SRCHLIB  
//SYSLIN DD DSNAME=&&OBJECT,DISP=(OLD,DELETE)  
// DD *  
ENTRY READIN  
REPLACE SEARCH(BINSRCH)  
INCLUDE SYSLMOD(TBLESRCH)  
NAME TBLESRCH(R)  
/*
```

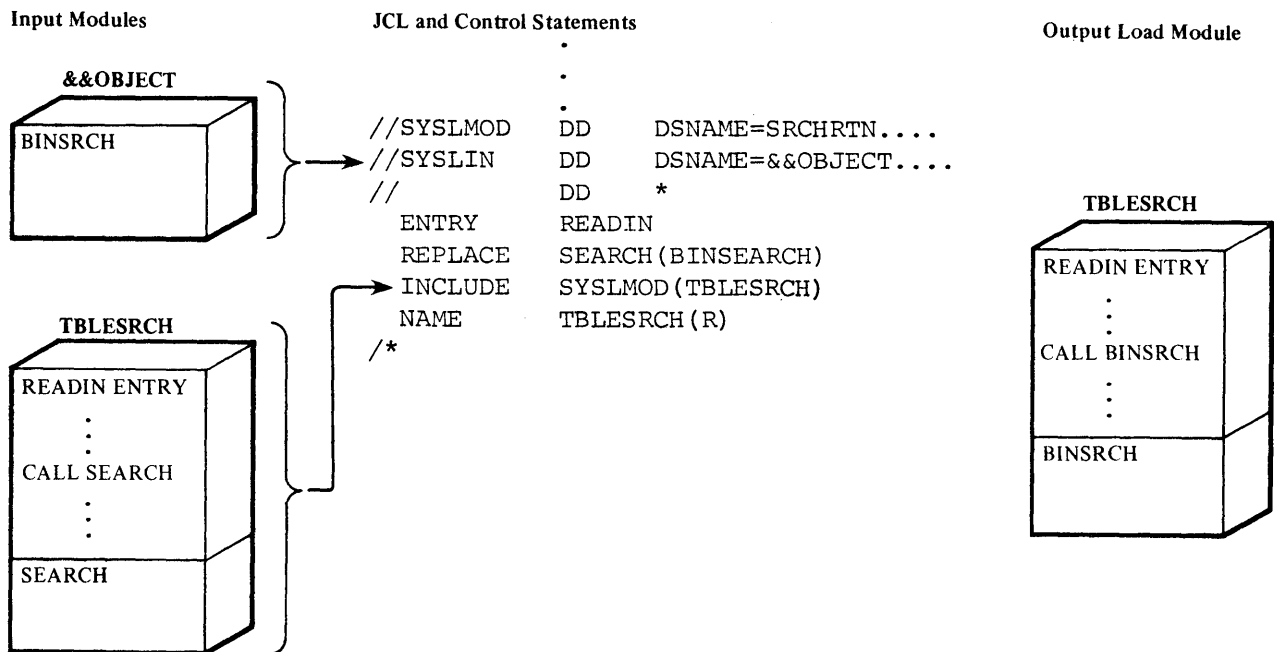


Figure 30. Replacing a Control Section with the REPLACE Control Statement

The output module contains BINSRCH instead of SEARCH; any references to SEARCH within the module refer to BINSRCH. Any external references to SEARCH from other modules will not be resolved to BINSRCH.

DELETING A CONTROL SECTION OR ENTRY NAME

The REPLACE statement can be used to delete a control section or an entry name. The REPLACE statement must immediately precede either the module that contains the control section or entry name to be deleted or the INCLUDE statement that specifies the module. Only one symbol appears on the REPLACE statement; the appropriate deletion is made depending on how the symbol is defined in the module.

If the symbol is a control section name, the entire control section is deleted. The control section name is deleted from the external symbol dictionary only if no address constants refer to the name from within the same input module. If an address constant does refer to it, the control section name is changed to an external record.

The preceding is also true of an entry name to be deleted. Any references to it from within the input module cause the entry name to be changed to an external reference.

These editor-supplied external references, unless resolved with other input modules, cause the automatic library call mechanism to attempt to resolve them. Also, the deletion of a control section or an entry name may cause external references from other input modules to be unresolved. Either condition can cause the output load module to be marked not executable.

If a deleted control section contains an unresolved external reference, the reference remains.

If a REPLACE statement, used to delete a CSECT, is the last control statement and there are external references to be resolved from SYSLIB, the delete request operates on the first module from SYSLIB and deletes it. The external reference remains unresolved.

Note: When a control section is deleted, any CSECT identification data associated with that control section is also deleted.

In the following example, control section CODER is to be deleted (Figure 31).

```

//SYSLMOD DD DSNAME=PVTLIB,DISP=OLD,UNIT=3380,
// VOLUME=SER=PVT002
//SYSLIN DD *
ENTRY START1
REPLACE CODER
INCLUDE SYSLMOD(CODEROUT)
NAME CODEROUT(R)
/*

```

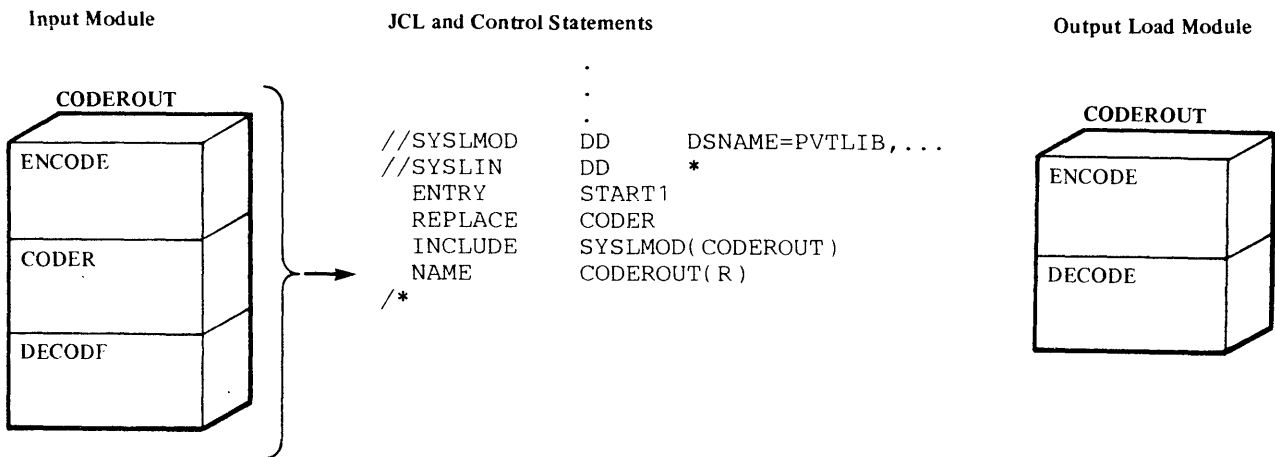


Figure 31. Deleting a Control Section

The control section CODER is deleted. If no address constants refer to CODER from other control sections in the module, the control section name is also deleted. If address constants refer to CODER, the name is retained as an external reference.

ORDERING CONTROL SECTIONS OR NAMED COMMON AREAS

The sequence of control sections or named common areas in an output load module can be specified by using the ORDER control statement.

Individual control sections or named common areas are arranged in the output load module according to the sequence in which they appear on the ORDER control statement. Multiple ORDER control statements can be used in a job step. The sequence of the ORDER statements determines the sequence of the control sections or named common areas in the load module.

Any control sections or named common areas that are not specified on ORDER statements appear last in the output load module. If a control section or named common area is changed by a CHANGE or REPLACE control statement, the new name must be used on the ORDER statement.



In the following example, ORDER statements are used to specify the sequence of five of the six control sections in an output load module. A REPLACE statement is used to replace the old control section, SESECTA, with the new control section, CSECTA, from the data set &&OBJECT, which was passed from a previous step. Assume that the control sections to be ordered are found in library member MAINROOT (Figure 32).

```

//SYSLMOD      DD      DSNAME=PVTLIB,DISP=OLD,
//              UNIT=3380,VOLUME=SER=PVT002
//SYSLIN       DD      DSNAME=&&OBJECT,DISP=(OLD,DELETE)
//              DD      *
//              ORDER  MAINEP,SEGMENT1,SEG2
//              REPLACE SESECTA(CSECTA)
//              ORDER  CSECTA,CSECTB
//              INCLUDE SYSLMOD(MAINROOT)
//              NAME    MAINROOT
/*

```

Input Modules



JCL and Control Statements

```

//          EXEC  PGM=HEWL
//          .
//          .
//          .
//SYSLMOD   DD      DSNAME=PVTLIB,...
//SYSLIN    DD      DSNAME=&&OBJECT,...
//          DD      *
//          ORDER  MAINEP(P),SEGMENT1,SEG2
//          REPLACE SESECTA(CSECTA)
//          ORDER  CSECTA,CSECTA,CSECTB(P)
//          INCLUDE SYSLMOD(MAINROOT)
//          NAME    MAINROOT
/*

```

Output Load Module




Figure 32. Ordering Control Sections

In the load module MAINROOT, the control sections MAINEP, SEGMENT1, SEG2, CSECTA, and CSECTB are rearranged in the output load module according to the sequence specified in the ORDER statements. A REPLACE statement is used to replace control section SESECTA with control section CSECTA from data set &&OBJECT, which was passed from a previous step. The ORDER statement refers to the new control section CSECTA. Control section LATESTP appears after the other control sections in the output load module, because it was not included in the ORDER statement operands.

ALIGNING CONTROL SECTIONS OR NAMED COMMON AREAS ON PAGE BOUNDARIES

A control section or named common area can be placed on a page boundary (to effect a lower paging rate and thus make more efficient use of real storage) by using either the ORDER statement or the PAGE statement.

The control section or common area to be aligned is named on either the PAGE statement or the ORDER statement with the P operand. Either the PAGE statement or the ORDER statement (with the P operand) causes the linkage editor to locate the starting address of the control section or common area on a page boundary within the load module.

In the following example, the control sections RAREUSE and MAINRT are aligned on page boundaries by PAGE and ORDER control statements. Control sections MAINRT, CSECTA, and SESECT1 are sequenced by the ORDER control statement. Assume that each control section, except for SESECT1 and RAREUSE, is 4K bytes in length (Figure 22).

```

//LKED      EXEC      PGM=HEWL,PARM='...'
           .
           .
//SYSLMOD   DD        DSNAME=OWNLIB,DISP=OLD,UNIT=3380,
//          VOLUME=SER=OWN002
//SYSLIN    DD        *
           PAGE      RAREUSE
           ORDER     MAINRT(P),CSECTA,SESECT1
           INCLUDE   SYSLMOD (MAINROOT)
           NAME      MAINROOT
/*

```


Input Module

JCL and Controls Statements

Output Load Module




Figure 33. Aligning Control Sections on Page Boundaries

The linkage editor places the control sections MAINRT and RAREUSE on page boundaries. Control sections MAINRT, CSECTA, and SESECT1 are sequenced as specified in the ORDER statement. RAREUSE, while placed on a page boundary, appears after the control sections specified in the ORDER statement because it was not included. The control section BOTTOM comes after RAREUSE because it appeared after RAREUSE in the input module.

CHAPTER 7. INVOKING THE LINKAGE EDITOR

The linkage editor can be invoked by a problem program at execution time through the use of one of the following macro instructions.

<u>[symbol]</u>	[LINK]	EP= <u>linkeditname</u> PARAM=(<u>optionlist</u> [, <u>ddname list</u>]), VL=1
<u>[symbol]</u>	[ATTACH]	EP= <u>linkeditname</u> PARAM=(<u>optionlist</u> [, <u>ddname list</u>]), VL=1
<u>[symbol]</u>	[LOAD]	EP= <u>linkeditname</u>
<u>[symbol]</u>	[XCTL]	EP= <u>linkeditname</u> PARAM=(<u>optionlist</u> [, <u>ddname list</u>]), VL=1

EP= linkeditname

specifies the symbolic name of the linkage editor. The entry point at which execution is to begin is determined by the control program (from the library directory entry). Any of the symbolic names that can be used as operands of the EXEC command's PGM parameter are acceptable as the "linkeditname".

PARAM=(optionlist[,ddname list])

specifies, as a sublist, address parameters to be passed from the problem program to the linkage editor. The first fullword in the address parameter list contains the address of the option and attribute list for the load module. The second fullword contains the address of the ddname list. If standard ddnames are to be used, this list may be omitted.

optionlist

specifies the address of a variable-length list containing the options and attributes. This address must be written even though no list is provided.

The option list must begin on a halfword boundary. The 2 high-order bytes contain a count of the number of bytes in the remainder of the list. If no options or attributes are specified, the count must be zero. The option list is free form, with each field separated by a comma. No blanks or zeros should appear in the list.

ddname list

specifies the address of a variable-length list containing alternative ddnames for the data sets used during linkage editor processing. If standard ddnames are used, this operand may be omitted.

The ddname list must begin on a halfword boundary. The 2 high-order bytes contain a count of the number of bytes in the remainder of the list. Each name of less than 8 bytes must be left justified and padded with blanks. If an alternate ddname is omitted from the list, the standard name will be assumed. If the name is omitted within the list, the 8-byte entry must contain binary zeros. Names can be omitted from the end by merely shortening the list.

The sequence of the 8-byte entries in the ddname list is as follows:

Entry	Alternate Name For:
1	SYSLIN
2	Member name (the name under which the output load module is stored in the SYSLMOD data set; this entry is used if the name is not specified on the SYSLMOD DD statement or if there is no NAME control statement)
3	SYSLMOD
4	SYSLIB
5	Not applicable
6	SYSPRINT
7	Not applicable
8	SYSUT1
9-11	Not applicable
12	SYSTEM

VL=1

specifies that the sign bit is to be set to 1 in the last fullword of the address parameter list.

When the linkage editor completes processing, a condition code is returned in register 15 (see Figure 16 on page 54 for a list of linkage editor return codes).

CHAPTER 8. INTERPRETING LINKAGE EDITOR OUTPUT

The linkage editor produces two types of output: a load module and diagnostic information. The principal output of the linkage editor is the output load module. The linkage editor always places this load module in a partitioned data set. In addition, the linkage editor issues diagnostic information. Error and/or warning messages, module disposition data, and optional diagnostic output are stored in the diagnostic output data set.

OUTPUT LOAD MODULE

The linkage editor produces one or more load modules (see "Load Module Format" on page 120) from the input processed. When more than one load module is produced, the process is called multiple load module processing.

Whether or not the linkage editor produces one or more load modules, the following apply:

- The load module is stored in a partitioned data set called the output module library.
- The load module must have an entry point; if the programmer has not assigned one, the linkage editor does.
- The output load module is assigned an authorization code.
- During processing, the linkage editor reserves and collects common areas, as specified in the source language program.
- During processing, the linkage editor accumulates total length and individual displacements for each pseudoregister (external dummy section).
- During processing, the linkage editor collects and records identification data in the CSECT identification (IDR) records.
- During the processing of a load module, the linkage editor deletes any private code (unnamed control section) having a length of zero and any identification data associated with it.
- The main entry point, each true alias, and each alternate entry point are assigned an addressing mode (AMODE).
- The output load module is assigned a residence mode (RMODE).

OUTPUT MODULE LIBRARY

The linkage editor stores every load module it produces in the output module library. This library is a partitioned data set that must be described by a DD statement with the name SYS1MOD. The data set name of the library is also specified on this DD statement. The data set can be either temporary (defined with a double ampersand), or permanent (defined with a single or no ampersand). If the data set name is either SYS1.LINKLIB or SYS1.SVCLIB, it would be advisable to re-IPL the system after linkage editor processing is complete. This ensures that the corresponding data extent block (DEB) is updated to reflect additional extents if secondary allocation of direct-access space was required.

Whether the data set is permanent or temporary, each module must be assigned a unique name, called the member name, to distinguish one load module from another. The output module can be assigned aliases if the programmer wants the module either identified by more than one name or entered for execution at several different points. Each member name and alias in a load module library must be unique. The library member name and aliases for each load module appear as separate entries in the library directory, along with the module attributes. (Some module attributes can be assigned on the EXEC statement for each linkage editor job step; see "Module Attributes" on page 38.)

Member Name

The member name of the output load module may be specified on the SYSLMOD DD statement, in a NAME statement, or both. If the member name is not specified, the default is TEMPNAME. If this default name has been previously assigned to a load module, using it again will cause a failure.

ASSIGNED ON SYSLMOD DD STATEMENT: If the member name is assigned on the SYSLMOD DD statement, the name is written in parentheses following the data set name of the library. For example:

```
//SYSLMOD DD DSNAME=MATHLIB(SQDEV),DISP=(NEW,KEEP),
// UNIT=3380,SPACE=(TRK,(100,10,1)),
// VOLUME=SER=LIB002
```

The member name SQDEV is assigned to the load module, which is placed in the new library named MATHLIB.

ASSIGNED ON NAME CONTROL STATEMENT: If the member name is not specified on the SYSLMOD DD statement, it may be assigned in a NAME control statement. For example:

```
//SYSLMOD DD DSNAME=MATHLIB,DISP=(NEW,KEEP),...
//SYSLIN DD DSNAME=&&OBJECT,DISP=(OLD,DELETE),...
// DD *
NAME SQDEV
/*
```

The member name SQDEV is assigned to the load module, which is placed in the library named MATHLIB.

ASSIGNED ON BOTH: If both the SYSLMOD DD statement and the NAME control statement specify a member name, the names should be identical. If the names are different, the name on the NAME control statement is used as the member name.

Note: If a "link-edit and go" sequence of job steps is performed and the program name in the EXEC statement of the "go" step contains a backward reference to the SYSLMOD DD statement in the "link-edit" step, the user must ensure that the member name specified in the SYSLMOD DD statement is valid and is not overridden by a NAME control statement.

An example of an error:

```
//LKED      EXEC   PGM=HEWL
           .
           .
//SYSLMOD   DD     DSNAME=&&LOADST(GO),DISP=(NEW,
//          PASS),...
//SYSLIN    DD     DSNAME=&&OBJECT,DISP=(OLD,DELETE),...
//          DD     *
//          NAME   READ
//          /*
//GO         EXEC   PGM=*.LKED.SYSLMOD
           .
           .
```

Remember, this example is incorrect!

The EXEC statement of the GO step specifies that the module to be executed is described in the LKED step in the SYSLMOD statement. The system tries to locate a member named GO; however, the output module was assigned the name READ.

REPLACING AN IDENTICALLY NAMED LIBRARY MEMBER: The output module can replace an identically named member in the library in either of two ways. The SYSLMOD DD statement names an existing data set, as follows:

```
//SYSLMOD   DD     DSNAME=MATHLIB(SQDEV),DISP=(OLD,
//          KEEP),...
```

Or, the NAME control statement specifies the replace function, as follows:

```
NAME       SQDEV(R)
```

In either case, the member named SQDEV is replaced with a new module of the same name.

Alias Names

An output module can be assigned a maximum of 16 aliases, specified with the ALIAS control statement. The aliases exist in addition to the member name of the output module. When a module is referred to by an alias, execution begins at the external name specified by the alias. If the name specified by the ALIAS statement is not an external symbol within the module, the main entry point is used.

For example, an output module is to be assigned two additional entry points, CODE1 and CODE2. In addition, because of a misunderstanding, calling modules have been written and tested using both ROUTONE and ROUT1 to refer to the output module. Rather than correct the calling modules, an alternate library member name (alias) is also assigned.

```

//SYSLMOD   DD      DSNAME=PVTLIB,DISP=OLD,UNIT=3380,
//          DD      VOLUME=SER=LIB001
//SYSLIN    DD      DSNAME=&&OBJECT,DISP=(OLD,DELETE)
//          DD      *
//          ALIAS   CODE1,CODE2,ROUTONE
//          NAME    ROUT1
/*

```

The names CODE1, CODE2, and ROUTONE appear in the library directory along with ROUT1, the member name. Because CODE1 and CODE2 are defined as external symbols within the output module, when these names are used, execution begins at these points. Control may be passed to the main entry point by using either the member name ROUT1 or the alias ROUTONE.

ENTRY POINT

Every load module must have a main entry point. The programmer may specify the entry point in one of two ways:

- On a linkage editor ENTRY control statement.
- On an Assembler language END statement, which is the last statement in the source program. The assembler produces an object module and an END statement for the module. The assembler-produced END statement contains an entry point only if the source language END statement contained one.

From its input, the linkage editor selects the entry point for the load module as follows:

1. From the first ENTRY control statement in the input.
2. If there is no ENTRY control statement in the input, from the first assembler-produced END statement that specifies an entry point.
3. If no ENTRY control statement or no assembler-produced END statement specifies an entry point, the first byte of the first control section of the load module is used as the entry point.

In general, the entry point should be explicitly specified, because it is not always possible to predict which control section will be first in the output module.

When a load module is reprocessed by the linkage editor, it has no END statement. Therefore, if the first byte of the first control section of the load module is not a suitable entry point, the entry point must be specified in one of two ways:

- Through an ENTRY control statement.
- Through the assembler-produced END statement of another input module, which is being processed for the first time. This object module must be the first such module to be processed by the linkage editor.

An entry point other than the main entry point may be specified with an ALIAS control statement. The symbol specified on the ALIAS statement must be defined as an external symbol in the load module. Any reference to that symbol causes execution of the module to begin at that point instead of at the main entry point.

In the following example, assume that CDCHECK, CODE1, and CODE2 are defined as external symbols in the output module:

```
//SYSLIN      DD      DSNAME=&&OBJECT,DISP=(OLD,DELETE)
//           DD      *
ENTRY CDCHECK
ALIAS CODE1,CODE2,ROUTONE
NAME ROUT1
/*
```

As a result of the preceding control statements, CDCHECK is the main entry point; CODE1 and CODE2 are alternate entry points. Any reference to ROUTONE or ROUT1 causes execution to begin at CDCHECK; any reference to CODE1 and CODE2 causes execution to begin at these points.

Authorization Code

Each load module link-edited is assigned an authorization code that determines whether or not the module is allowed to use restricted system services and resources. A nonzero code allows the module to use restricted services and resources; a zero code disallows that usage. The authorization code becomes part of the directory entry for the module in the library containing the module.

Residence and Addressing Modes

Each entry in the library directory for the output load module (one for the main entry point and one for each true alias or alternate entry point) contains an indication of the residence mode for the load module and an indication of the addressing mode for that entry point. The entries for true aliases and alternate entry points also contain an indication of the addressing mode for the main entry point.

RESERVING STORAGE IN THE OUTPUT LOAD MODULE

In FORTRAN, Assembler language, and PL/I, the programmer can create control sections that reserve virtual storage areas that contain no data or instructions. These control sections are called "common" or "static external" areas, and are produced in the object modules by the language translators. These common areas are used, for example, as communication regions for different parts of a program or to reserve virtual storage areas for data supplied at execution time. These common areas are either named or unnamed (blank).

COLLECTION OF COMMON AREAS: During processing, the linkage editor collects common areas. That is, if two or more blank common areas are found in the input, the largest blank common area is used in the output module; all references to a blank common area refer to the one retained. If two or more named common areas have the same name, the largest of the identically named common areas is used in the output module; all references to the named common areas refer to the one area retained.

IDENTICALLY NAMED COMMON AREAS AND CONTROL SECTIONS: If a control section (as is generated from a BLOCK DATA subprogram in FORTRAN, for example) and a named common area have the same name, the length of the control section must be greater than or equal to the length of the named common area. If the control section is smaller in length than the named common area, a diagnostic message is issued. The control section is regarded as the largest of the common areas processed with that name. All subsequent control sections and/or common areas with the same name are ignored.

PROCESSING PSEUDOREGISTERS

In PL/I, programmers can use pseudoregisters to define storage that will not be reserved in the load module but can be allocated dynamically during execution. The external dummy sections generated by Assembler H Version 2 correspond to the pseudoregisters of PL/I.

The linkage editor accumulates the total length of all pseudoregisters in the input and records the displacement of each. If two or more pseudoregisters have the same name, the one with the longest length and the most restrictive alignment will be retained. All other pseudoregisters with the same name will be ignored; all references to the identically named pseudoregisters will refer to the one retained.

MULTIPLE LOAD MODULE PROCESSING

The linkage editor can produce more than one load module in a single job step. A NAME control statement in the input stream is used as a delimiter for input to a load module. If additional input modules follow the NAME statement in the input stream, they are used in the formation of the next load module.

Each load module that is formed has a unique name and is placed in the same library as a separate member. When processing multiple load modules in a single job step, the options and attributes specified in the EXEC statement for that job step apply to all load modules created. If the linkage editor terminates abnormally during processing of any of the output modules, neither that module nor any of the modules yet to be processed in the job step is processed or placed in the library. Load modules processed before abnormal termination have already been placed in the library.

In the following example, two load modules are produced in one linkage editor job step:

```

//LKED          EXEC  PGM=HEWL,PARM='MAP,LIST'
                .
                .
//SYSLMOD       DD    DSNAME=PAYROLL(OVERTIME),DISP=OLD,
//              UNIT=3380,VOLUME=SER=LIB002
                .
                .
//MODTWO        DD    DSNAME=&&OBJECT,DISP=(OLD,DELETE)
//SYSLIN        DD    DSNAME=&&OBJECT(A),DISP=(OLD,DELETE)
//              DD    *
    ENTRY       INIT
    NAME        OVERTIME
    INCLUDE     MODTWO(B)
    ENTRY       HSKEEP
    NAME        VACATION
/*

```

The first load module is produced from the object module in the data set defined on the SYSLIN DD statement. The main entry point is INIT and the member name is OVERTIME.

The second load module is produced from the object module specified by the INCLUDE statement. The main entry point is HSKEEP and the member name is VACATION.

If an INCLUDE statement specifies a member name that is different from the member name on the DD statement, the member specified on the DD statement must exist even though it is not to be included.

Both load modules are placed in the library PAYROLL, defined on the SYSLMOD statement.

The parameters on the EXEC card specify that a module map and a control statement listing are produced for each load module. The map and listing are discussed in detail in the next section.

DIAGNOSTIC OUTPUT

Diagnostic information is stored in the diagnostic output data set, which must be defined by a DD statement with the name SYSPRINT. This output is a collection of messages generated by the linkage editor, as well as any optional output requested by the programmer.

DIAGNOSTIC MESSAGES

The linkage editor generates two types of messages: module disposition messages and error/warning messages. Descriptions of the error/warning messages will be found in System Messages.

Module Disposition Messages

Module disposition messages of several types are printed for each load module produced. The first message indicates the options and attributes specified for each module. Invalid options or attributes are replaced by INVALID in the output. Messages are also generated to inform the programmer that incompatible attributes have been specified.

Disposition messages also describe the handling of the load module. These messages are preceded by several asterisks, and are:

- member name NOW ADDED TO DATA SET.
- member name NOW REPLACED IN DATA SET.
- member name DOES NOT EXIST BUT HAS BEEN ADDED TO THE DATA SET.

The replacement function was specified, but the member did not exist in the data set; the module is added to the data set using the member name given.

- alias name IS AN ALIAS FOR THIS MEMBER.
- MODULE HAS BEEN MARKED NOT EXECUTABLE.

In addition, module disposition messages are used when the reenterable (RENT), reusable (REUS), and/or refreshable (REFR) linkage editor options have been specified for the module. When one or more of these module attributes have been indicated, a message informs the user what attribute(s) have been assigned to the module. This message indicates whether the load module has been marked reenterable or not reenterable, reusable or not reusable, refreshable or not refreshable,

depending on the option or options used. (See "Reusability Attributes" on page 40 and "Refreshable Attribute" on page 41 for more information on these options.)

The message consists of several asterisks and MODULE HAS BEEN MARKED, followed by the attribute(s) assigned as a result of the linkage editor options specified. The programmer, of course, is responsible for verifying that the module actually is reenterable, reusable, and/or refreshable. The following messages are examples of some possible combinations:

- MODULE HAS BEEN MARKED REFRESHABLE.
- MODULE HAS BEEN MARKED NOT REFRESHABLE.
- MODULE HAS BEEN MARKED REUSABLE AND NOT REFRESHABLE.
- MODULE HAS BEEN MARKED REUSABLE AND REFRESHABLE.

When an error causes the linkage editor to mark a module not executable, only the MODULE HAS BEEN MARKED NOT EXECUTABLE message appears; no attribute messages are generated.

Error/Warning Messages

Certain conditions that are present when a module is being processed can cause an error or warning message to be printed. These messages contain a message code and message text. If an error is encountered during processing, the message code for that error is printed with the applicable symbol or record in error. After processing is completed, the diagnostic message associated with that code is printed. The error warning messages have the following format:

IEW0mms message text

where:

IEW0 indicates a linkage editor message

mm is the message number

s is the severity code, and may be one of the following values:

- 1 Indicates a condition that may cause an error during execution of the output module. A module map or cross-reference table is produced if specified by the programmer. The output module is marked executable.
- 2 Indicates an error that could make execution of the output module impossible. Processing continues. When possible, a module map or a cross-reference table is produced if specified by the programmer. The output module is marked not executable, unless the LET option is specified on the EXEC statement.
- 3 Indicates an error that will make execution of the output module impossible. Processing continues. When possible, a module map or a cross-reference table is produced if specified by the programmer. The output module is marked not executable.
- 4 Indicates an error condition from which no recovery is possible. Processing terminates. The only output is diagnostic messages.

Note: A special severity code of zero is generated for each control statement printed as a result of the LIST option. Severity zero does not indicate an error warning condition.

The highest severity code encountered during processing is multiplied by 4 to create a return code that is placed in register 15 at the end of processing. This return code can be tested to determine whether or not processing is to continue (see "EXEC Statement—Return Code" on page 54).

message text contains combinations of the following:

- The message classification (either error or warning)
- Cause of error
- Identification of the symbol, segment number (when in overlay), or input item to which the message applies
- Instructions to the programmer
- Action taken by the linkage editor

Optionally, error/warning messages can be sent to a separate output data set, which is defined by specifying TERM in the PARM field of the EXEC statement and including a SYSTEM DD statement. This separate SYSTEM data set consists of only numbered error/warning messages. It supplements the SYSPRINT output data set, which can also include module disposition messages and optional diagnostic output. When SYSTEM is used, the numbered error/warning messages appear in both data sets.

System Messages contains a complete list of error/warning messages.

Sample Diagnostic Output

Figure 34 on page 118 shows the format of the diagnostic output for the linkage editor. No optional output was requested other than the list of control statements.

The letters indicate the disposition and error/warning messages as follows:

- A Is a module disposition message that lists the options and attributes specified. Additional information is printed indicating the variable and default options used.
- B Is a list of control statements used (IEW0000) and the message codes (IEW0201 and IEW0461) for error/warning conditions discovered during processing. For error/warning message codes, the symbol in error, if necessary, is also listed (CCCCCCC and BASEDUMP).
- C Is a module disposition message (XXXX) that indicates that the output module (BBBBBBBB) has been added to the output module data set.
- D Is the diagnostic message directory that contains the text of the error codes listed in item B.

OPTIONAL OUTPUT

In addition to error/warning and disposition messages, the linkage editor can produce diagnostic output as requested by the programmer. This optional output includes a control statement listing, a module map, and a cross-reference table.

```

A      F64-LEVEL LINKAGE EDITOR OPTIONS SPECIFIED LET,NCAL,XREF,OVLY,LIST
      DEFAULT OPTIONS(S) USED - SIZE=(65536,6144)
B      IEW0000 NAME BBBBBBBB
      IEW0201
      IEW0461 CCCCCCCC
      IEW0461 BASEDUMP
C      *****BBBBBBBB NOW ADDED TO DATA SET
      DIAGNOSTIC MESSAGE DIRECTORY
D      IEW0201 WARNING - OVERLAY STRUCTURE CONTAINS ONLY ONE SEGMENT -- OVERLY OPTION
      CANCELED.
      IEW0461 WARNING - SYMBOL PRINTED IS AN UNRESOLVED EXTERNAL REFERENCE, NCAL WAS
      SPECIFIED.

```

Figure 34. Diagnostic Messages Issued by the Linkage Editor

Control Statement Listing

If the LIST option is specified on the EXEC statement, a listing of all linkage editor control statements is produced. For each control statement, the listing contains a special message code, IEW0000, followed by the control statement. Item B in Figure 34 contains an example of a control statement listing.

Module Map

If the MAP option is specified on the EXEC statement, a module map of the output load module is produced. The module map shows all control sections in the output module and all entry names in each control section. Named common areas are listed as control sections.

For each control section, the module map indicates its origin (relative to zero) and length in bytes (in hexadecimal notation). For each entry name in each control section, the module map indicates the location at which the name is defined. These locations are also relative to zero.

If the module is not in an overlay structure, the control sections are arranged in ascending order according to their origins. An entry name is listed with the control section in which it is defined.

If the module is an overlay structure, the control sections are arranged by segment. The segments are listed as they appear in the overlay structure, top to bottom, left to right, and region by region. Within each segment, the control sections and their corresponding entry names are listed in ascending order according to their assigned origins. The number of the segment in which they appear is also listed.

In any module map, the following are identified by a dollar sign:

- Blank common area
- Private code (unnamed control section)
- For overlay programs, the segment table and each entry table

When the load module processed by the linkage editor does not have an origin of zero, the linkage editor generates a one-byte private code (unnamed control section) as the first text record. This private code is deleted in any subsequent reprocessing of the load module by the linkage editor.

Each control section that is obtained from a call library during automatic library call is identified by an asterisk after the control section name.

At the end of the module map is the entry address, that is, the relative address of the main entry point. The entry address is followed by the total length of the module in bytes; in the case of an overlay module, the length is that of the longest path. Pseudoregisters, if used, also appear at the end of the module map; the name, length, and displacement of each pseudoregister are given.

Figure 35 contains a module map with five control sections. There are two named control sections (COBSUB and MAINMOD), one unnamed control section (designated by \$PRIVATE), and two control sections obtained from a call library (ILBODSPO and ILBOSTPO). In addition, two entry names are defined: SUB1 in the unnamed control section and ILBOSTP1 in control section ILBOSTPO.

CONTROL SECTION			ENTRY							
NAME	ORIGIN	LENGTH	NAME	LOCATION	NAME	LOCATION	NAME	LOCATION	NAME	LOCATION
COBSUB	00	33A								
\$PRIVATE	340	EF	SUB1	340						
MAINMOD	430	166								
ILBODSPO*	598	5E2								
ILBOSTPO*	B80	35	ILBOSTP1	B96						
ENTRY ADDRESS	430									
TOTAL LENGTH	BB8									

****GO DOES NOT EXIST BUT HAS BEEN ADDED TO DATA SET

Figure 35. Module Map

Cross-Reference Table

If the XREF option is specified on the EXEC statement, a cross-reference table is produced. The cross-reference table consists of a module map and a list of cross-references for each control section. Each address constant that refers to a symbol defined in another control section is listed with its assigned location, the symbol referred to, and the name of the control section in which the symbol is defined. When control sections are compiled together, and simple address constants are used to refer from one control section to another (instead of using external symbols and entry names), the control section name is listed as the symbol referred to.

For overlay programs, this information is provided for each segment; in addition, the number of the segment in which the symbol is defined, is provided.

If a symbol is unresolved after processing by the linkage editor, it is identified by \$UNRESOLVED in the list. However, if an unresolved symbol is marked by the never-call function (as specified on a LIBRARY control statement), it is identified by \$NEVER-CALL. If an unresolved symbol is a weak external reference, it is identified by \$UNRESOLVED(W).

Figure 36 on page 120 contains a cross-reference table for the same program whose module map is shown in Figure 35. All the information from the module map is present, plus a list of cross-references for each control section.

CROSS-REFERENCE TABLE

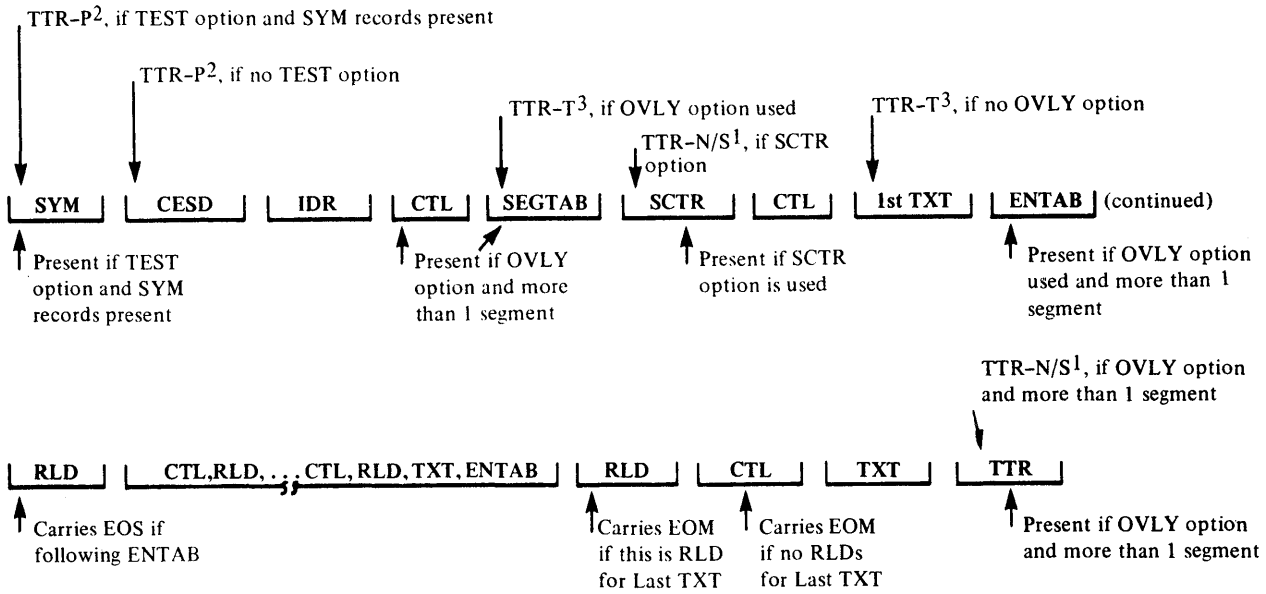
CONTROL SECTION			ENTRY		NAME		LOCATION		NAME		LOCATION	
NAME	ORIGIN	LENGTH	NAME	LOCATION	NAME	LOCATION	NAME	LOCATION	NAME	LOCATION	NAME	LOCATION
COBSUB	00	33A										
\$PRIVATE	340	EF										
MAINMOD	430	166	SUB1	340								
ILBODSPO*	598	5E2										
ILBOSTPO*	B80	35	ILBOSTP1	B96								
LOCATION	REFERS TO SYMBOL	IN CONTROL SECTION	LOCATION	REFERS TO SYMBOL	IN CONTROL SECTION							
250	ILBOSTPO	ILBOSTPO	254	ILBODSPO	ILBODSPO							
258	ILBOSTP1	ILBOSTPO	450	SUB1	SUB1							
478	COBSUB	COBSUB										
ENTRY ADDRESS	430											
TOTAL LENGTH	BB8											

Figure 36. Cross-Reference Table

LOAD MODULE FORMAT

The format of a load module built by the linkage editor is shown in Figure 37 on page 121.

In writing the output load module to the SYSLMOD data set, the linkage editor does not use the track overflow feature. When moving or copying load modules, the track overflow feature must not be used on the target data set, as errors may occur in fetching the load modules for execution.



¹TTR-N/S: TTR of the note list or scatter/translation table. Used for modules in scatter load format or overlay structure only.

²TTR-P: TTR of the first block of the named member (load module).

³TTR-T: TTR of the first block of text.

Figure 37. Load Module Format

APPENDIX A. SAMPLE LINKAGE EDITOR PROGRAMS

This appendix contains sample linkage editor programs. The material presented for each program includes a description of the program, the job control language necessary for the linkage editor job step, linkage editor control statements (if any), and the linkage editor output. The sample programs are:

- Link-editing a COBOL and a FORTRAN object module (COBFORT)
- Replacing one control section with another by using the REPLACE statement (RPLACJOB)
- Creating a multiple-region overlay program (REGNOVLY)
- Placing the control statements for the multiple region overlay program in a partitioned data set, and using them (PARTDS)

The output for each program includes a cross-reference table, a module map, a control statement listing, and diagnostic messages, if any.

SAMPLE PROGRAM COBFORT

Sample program COBFORT link-edits a COBOL object module and a FORTRAN object module to form one load module. The source programs were compiled in two steps previous to the linkage editor job step, and the output from each compilation was placed in data set &&OBJMOD.

Job Control Language

The job control language for the linkage editor job step of this sample program is:

```
//LKED      EXEC  PGM=HEWL,PARM='XREF'  
//SYSUT1   DD    DSNAME=&&UT1,UNIT=SYSDA,SPACE=(TRK,  
//          (100,10))  
//SYSLIB   DD    DSNAME=SYS1.COBLIB,DISP=SHR  
//          DD    DSNAME=SYS1.FORTLIB,DISP=SHR  
//SYSLMOD  DD    DSNAME=&&LOADMD(GO),UNIT=SYSDA,  
//          DISP=(NEW,PASS),SPACE=(TRK,  
//          (100,10,1))  
//SYSPRINT DD    SYSOUT=A  
//SYSLIN   DD    DSNAME=&&OBJMOD,DISP=(OLD,DELETE)  
/x
```

Statement	Explanation
EXEC	Causes the execution of the linkage editor. The PARM field option requests a cross-reference table and a module map to be produced on the diagnostic output data set.
SYSUT1	Defines a temporary direct access data set to be used as the intermediate data set.

SYSLIB Defines the automatic call library; the call libraries for COBOL and FORTRAN are concatenated; both are used to resolve external references.

SYSLMOD Defines a temporary data set to be used as the output module library; the load module is assigned a member name of GO, and is passed to a subsequent step for execution.

SYSPRINT Defines the diagnostic output data set, which is assigned to output class A.

SYSLIN Defines the primary input data set, &&OBJMOD, which contains both input object modules; this data set was passed from a previous job step and is to be deleted at the end of this job step.

Linkage Editor Output

Figure 38 on page 124 shows the linkage editor output for COBFORT. The listing header indicates the options specified (XREF), and the SIZE option values in decimal (196608 for value1 and 65536 for value2). Because XREF is specified, the heading CROSS REFERENCE TABLE precedes the rest of the output.

Figure 38 also shows the module map for COBFORT. IPCT30 and TX652F are the names of the input control sections. The rest of the control sections are either from the COBOL automatic call library or from the FORTRAN automatic call library. (They can be distinguished by the initial three letters; ILB indicates a COBOL control section, IHC a FORTRAN control section.) The origin and length (in hexadecimal) of each control section follow the name.

To the right of each control section is a list of the entry names defined in each control section. The location (in hexadecimal) of each entry name is also given. For example, in control section IHCCOMH2 (the asterisk is not a part of the name; it indicates that the control section is from the automatic call library), entry name SEQDASD is defined at location 154A.

Figure 38 shows the cross-reference table for COBFORT. The table contains the location of any address constant that refers to a symbol defined in another control section. The symbol the address constant refers to is also listed, along with the control section in which the symbol is defined. For example, at location 1F0 in control section IPCT30 (determined by using the module map; 1F0 falls between origin 00 and origin 360), an address constant refers to symbol IHDFDISP, defined in control section IHDFDISP.

The entry address is 00 and the total length of the load module is 4AE8. Note that the length of the module is rounded up to a doubleword boundary.

The disposition message at the end of the output in Figure 38 indicates that the load module GO has been added to the output module library. The library did not contain any other module with that name. The four asterisks identify the message.

SAMPLE PROGRAM RPLACJOB

Sample program RPLACJOB shows the use of the REPLACE statement to replace one control section with another. The source program for the new control section (NEWMOD) is processed in a previous job step and passed to the linkage editor job step. The control section (SUBONE) to be replaced is in an existing load module. Figure 39 on page 125 shows the linkage editor output for the job step that created this load module. Note that the entry address is F0, which is the location of the entry point MAINMOD (specified on the ENTRY control statement).

Module Map

P64-LEVEL LINKAGE EDITOR OPTIONS SPECIFIED XREF
 DEFAULT OPTIONS(S) USED - SIZE=(196608,65536)

CROSS REFERENCE TABLE

CONTROL SECTION			ENTRY							
NAME	ORIGIN	LENGTH	NAME	LOCATION	NAME	LOCATION	NAME	LOCATION	NAME	LOCATION
IPCT30	00	360								
TX652F	360	1E0								
IHCPCOMH*	540	CD9								
IHCCOMH2*	1220	434	IBCOM#	540	FDIOCS#	5FC	INTSWTCH	11FE		
IHDFDISP*	1658	626	SEQDASD	154A						
IHCPCVTH*	1C80	119D								
			ADCON#	1C80	FCVAOUTP	1D2A	FCVLOUTP	1DBA	FCVZOUTP	1FOA
			FCVIOUTP	22B8	FCVEOUTP	27BA	FCVCOUTP	29D4	INT6SWCH	2CBB
IHCFINTH*	2E20	39E								
IHCFIOSH*	31C0	100E	ARITH#	2E20	ADJSWTCH	30D8				
IHCUIOPT *	41D0	8	FIOCS#	31C0						
IHCTRCH *	41D8	2D4								
IHCUIATBL*	44B0	638	IHCERRM	41D8						

Cross-Reference Table

LOCATION	REFERS TO SYMBOL	IN CONTROL SECTION	LOCATION	REFERS TO SYMBOL	IN CONTROL SECTION
1F0	IHDFDISP	IHDFDISP	1F4	TX652F	TX652F
410	IBCOM#	IHCPCOMH	5FC	SEQDASD	IHCCOMH2
1108	ADCON#	IHCPCVTH	1100	FIOCS#	IHCFIOSH
110C	ARITH#	IHCFINTH	112C	ADJSWTCH	IHCFINTH
1128	IHCUIOPT	IHCUIOPT	1110	FCVEOUTP	IHCPCVTH
1114	FCVLOUTP	IHCPCVTH	1118	FCVIOUTP	IHCPCVTH
111C	FCVCOUTP	IHCPCVTH	1120	FCVAOUTP	IHCPCVTH
1124	FCVZOUTP	IHCPCVTH	10E0	IHCCOMH2	IHCCOMH2
10E4	IHCERRM	IHCTRCH	14A9	IHCPCOMH	IHCPCOMH
14AC	IHCPCOMH	IHCPCOMH	1268	IHCERRM	IHCTRCH
1264	IBCOM#	IHCPCOMH	2C7C	IBCOM#	IHCPCOMH
2C78	IHCERRM	IHCTRCH	311C	IBCOM#	IHCPCOMH
3120	INTSWTCH	IHCPCOMH	30D4	INT6SWCH	IHCPCVTH
30D0	IHCUIOPT	IHCUIOPT	3128	ADCON#	IHCPCVTH
3124	FIOCS#	IHCFIOSH	32F8	IHCERRM	IHCTRCH
3FP8	IHCUIATBL	IHCUIATBL	4004	IBCOM#	IHCPCOMH
43D0	IBCOM#	IHCPCOMH	43D4	ADCON#	IHCPCVTH
43D8	FIOCS#	IHCFIOSH			

ENTRY ADDRESS 00
 TOTAL LENGTH 4AEB

****GO DOES NOT EXIST BUT HAS BEEN ADDED TO DATA SET
 AUTHORIZATION CODE IS 0.

Figure 38. Linkage Editor Output for Sample Program COBFORT

F64-LEVEL LINKAGE EDITOR OPTIONS SPECIFIED XREF,LIST
 DEFAULT OPTION(S) USED - SIZE=(196608,65536)
 LEW0000 ENTRY MAINMOD

CROSS REFERENCE TABLE

CONTROL SECTION			ENTRY					
NAME	ORIGIN	LENGTH	NAME	LOCATION	NAME	LOCATION	NAME	LOCATION
SUBONE	00	EF						
MAINMOD	F0	146	SUB1	00				

LOCATION	REFERS TO SYMBOL	IN CONTROL SECTION	LOCATION	REFERS TO SYMBOL	IN CONTROL SECTION
11C	SUBONE	SUBONE			
ENTRY ADDRESS	F0				
TOTAL LENGTH	238				

****GO DOES NOT EXIST BUT HAS BEEN ADDED TO DATA SET
 AUTHORIZATION CODE IS 0.

Figure 39. Linkage Editor Output for Job Step that Created SUBONE

Job Control Language

The job control language for the replacement job step of this sample program is shown below.

```

//LKED      EXEC   PGM=HEWL, PARM='XREF,LIST'
//SYSUT1    DD     UNIT=SYSDA, SPACE=(TRK,(100,10))
//INPUTX    DD     DSNAME=LOADLIB, DISP=OLD, UNIT=SYSDA,
//           VOL=SER=SCRATCH
//SYSLMOD   DD     DSNAME=LOADLIB(GO), DISP=OLD, UNIT=SYSDA,
//           VOL=SER=SCRATCH
//SYSPRINT  DD     SYSOUT=A
//SYSLIN    DD     DSNAME=&&OBJMOD, DISP=(OLD,DELETE),
//           UNIT=SYSDA
//           DD     *
Linkage Editor Control Statements
/*

```

Statement	Explanation
EXEC	Causes the execution of the linkage editor. The PARM field options request a cross-reference table and a module map (XREF), and a control statement listing (LIST) to be produced on the diagnostic output data set.
SYSUT1	Defines a temporary direct access data set to be used as the intermediate data set.
INPUTX	Defines a permanent data set, used later as additional linkage editor input.
SYSLMOD	Defines a permanent data set to be used as the output module library. Note that it is the same data set that was described on the INPUTX DD statement. The output load module is added to the data set, under the member name GO.
SYSPRINT	Defines the diagnostic output data set, which is assigned to output class A.
SYSLIN	Defines the primary input data set, &&OBJMOD, which contains the object module for the replacement control section. This data set is temporary and was passed from a previous job step; it is to be deleted at the end of this job. This statement also concatenates the input stream to the primary input data set. The input stream contains linkage editor control statements that must be followed by a /* statement.

Figure 40. Job Control Statements for RPLACJOB

LINKAGE EDITOR CONTROL STATEMENTS

The input stream contains the linkage editor control statements that are necessary for the replacement of SUBONE with NEWMOD. The control statements are shown below:

```

ENTRY    MAINMOD
REPLACE  SUBONE(NEWMOD)
INCLUDE  INPUTX(GO)

```

Statement	Explanation
ENTRY	Specifies that the entry point is to be MAINMOD.
REPLACE	Specifies that control section SUBONE in the module that follows the REPLACE statement is to be replaced by control section NEWMOD.
INCLUDE	Specifies additional input: member GO of the data set described on the INPUTX DD statement. This library member contains the control section to be replaced. Because this member name is identical to that specified on the SYSLMOD DD statement, the output load module replaces the existing library member.

Figure 41. Linkage Editor Control Statements for RPLACJOB

Linkage Editor Output

Figure 42 shows the linkage editor output for sample program RPLACJOB. The listing header indicates the options specified (XREF and LIST), and the SIZE option values used (196608 for value1 and 65536 for value2).

```
F64-LEVEL LINKAGE EDITOR OPTIONS SPECIFIED XREF,LIST
DEFAULT OPTION(S) USED - SIZE=(196608,65536)
IEW0000      ENTRY MAINMOD
IEW0000      REPLACE SUBONE(NEWMOD)
IEW0000      INCLUDE INPUTX(GO)
```

CROSS REFERENCE TABLE

CONTROL SECTION			ENTRY					
NAME	ORIGIN	LENGTH	NAME	LOCATION	NAME	LOCATION	NAME	LOCATION
NEWMOD	00	F1						
MAINMOD	F8	146						

LOCATION	REFERS TO SYMBOL	IN CONTROL SECTION	LOCATION	REFERS TO SYMBOL	IN CONTROL SECTION
124	NEWMOD	NEWMOD			

ENTRY ADDRESS F8
TOTAL LENGTH 240
****GO NOW REPLACED IN DATA SET
AUTHORIZATION CODE IS 0.

Figure 42. Linkage Editor Output for Sample Program RPLACJOB

Because the LIST option is specified, a control statement listing is produced. Each control statement is preceded by a special message number, IEW0000. Because XREF is specified, the heading CROSS REFERENCE TABLE precedes the rest of the output.

The module map shows that control section NEWMOD is now part of the load module, and that control section SUBONE has been deleted. The new entry address is F8, because NEWMOD is longer than SUBONE. The total length of the load module is 240 bytes.

The cross-reference table indicates that at location 124 in MAINMOD, an address constant refers to symbol NEWMOD, defined in control section NEWMOD. Note that before the replacement occurred, the address constant in MAINMOD referred to SUBONE, defined in control section SUBONE (Figure 39 on page 125). When the REPLACE statement is used to replace a control section, references to the old control section from within the same input module are also changed.

The disposition message indicates that the output load module (GO) has been added to the output module library.

SAMPLE PROGRAM REGNOVLY

Sample program REGNOVLY creates a multiple-region overlay structure. The structure produced is shown in Figure 43 on page 128. In this program, some of the references between control sections are:

CSA to CSE

CSB to CSE

CSB to CSD

CSD to CSC

The reference from CSB to CSE is a valid exclusive call, because there is a reference to CSE in the segment common to

REGION 1

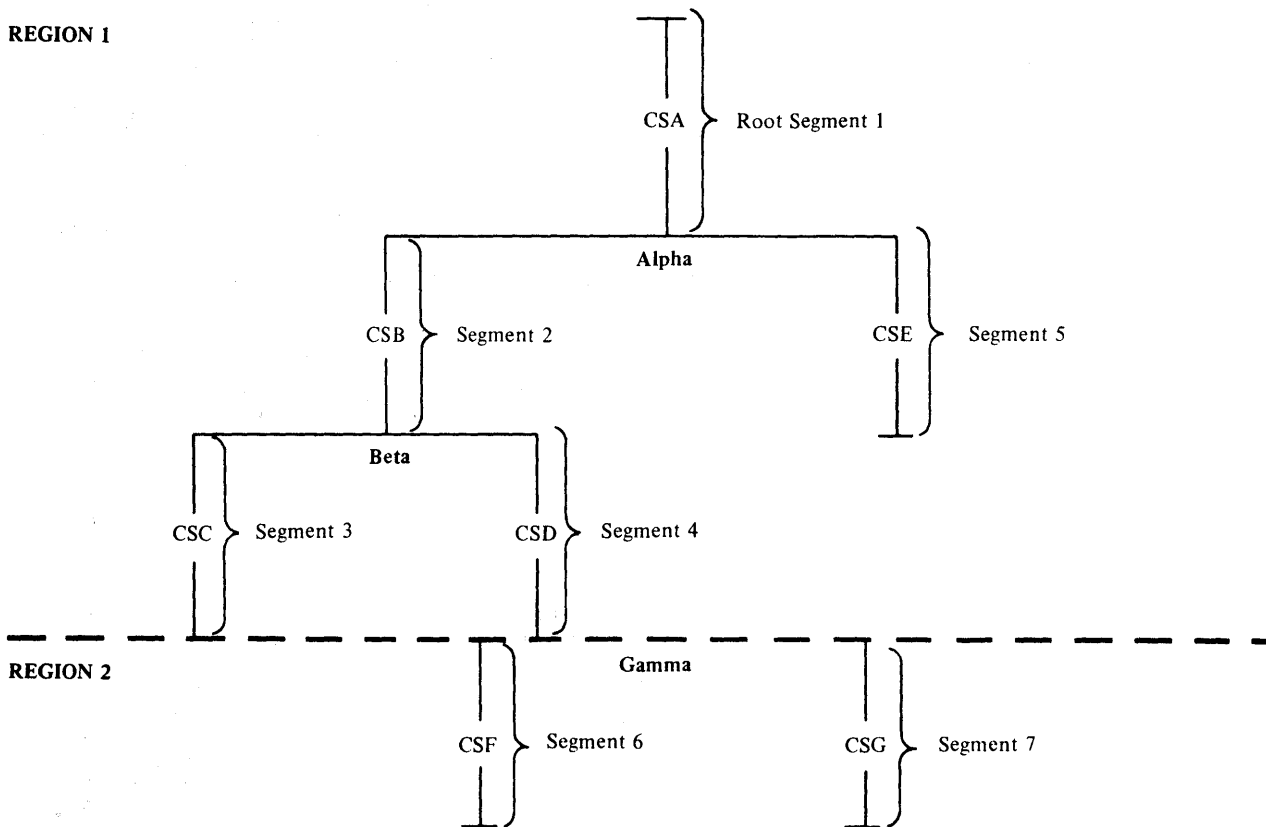


Figure 43. Overlay Tree for Multiple-Region Sample Program REGNOVLY

both CSB and CSE; the reference from CSD to CSC is invalid, because there is no reference to CSC in the common segment.

The source programs for all the control sections were compiled in previous job steps. All the object modules were placed in the same data set, which was passed to the linkage editor job step.

Job Control Language

The job control language for the linkage editor job step of this sample program is shown below.

```
//LKED      EXEC   PGM=HEWL,PARM='XREF,LIST,OVLY,LET'  
//SYSUT1   DD     DSNAME=&&UT1,UNIT=SYSDA,SPACE=(TRK,  
//          (100,10))  
//SYSLIB   DD     DSNAME=SYS1.COBLIB,DISP=SHR  
//SYSLMOD  DD     DSNAME=&&OVLYJB(GO),UNIT=SYSDA,  
//          DISP=(NEW,PASS),SPACE=(TRK,(100,10,1))  
//SYSPRINT DD     SYSOUT=A  
//SYSLIN   DD     DSNAME=&&OBJMOD,DISP=(OLD,DELETE)  
//          DD     *  
Linkage Editor Control statements  
/*
```

Statement	Explanation
EXEC	Causes the execution of the linkage editor. The PARM field options request a cross-reference table and a module map (XREF), and a control statement listing (LIST) to be produced on the diagnostic output data set. The module is to be assigned the overlay attribute (OVLY), and marked executable in spite of severity 2 errors (LET). The LET option is specified to permit testing of the output module, even though an invalid exclusive call is present. The XCAL option allows only valid exclusive calls.
SYSUT1	Defines a temporary direct access data set to be used as the intermediate data set.
SYSLIB	Defines the automatic call library (SYS1.COBLIB) to be used to resolve external references. All control sections from this library are placed in the root segment; they remain there unless they are repositioned.
SYSLMOD	Defines a temporary data set to be used as the output module library; the load module is assigned the member name GO and is passed to a subsequent step for execution.
SYSPRINT	Defines the diagnostic output data set, which is assigned to output class A.
SYSLIN	Defines the primary input data set, &&OBJMOD, which contains the object modules for the overlay structure. This data set is temporary and was passed from a previous job step; it is to be deleted at the end of this job. This statement also concatenates the input stream to the primary input data set. The input stream contains linkage editor control statements, which must be delimited by a /* statement.

Figure 44. Job Control Statements for REGNOVLY

Linkage Editor Control Statements

The input stream contains the linkage editor control statements that structure the overlay program. The control statements are:

```

INSERT CSA
ENTRY CSA
OVERLAY ALPHA
INSERT CSB
OVERLAY BETA
INSERT CSC
OVERLAY BETA
INSERT CSD
OVERLAY ALPHA
INSERT CSE
OVERLAY GAMMA(REGION)
INSERT CSF
OVERLAY GAMMA
INSERT CSG

```


Linkage Editor Output

Figure 45 shows the linkage editor output for sample program REGNOVLY. The list header indicates the options specified and the SIZE option values used.

P64-LEVEL LINKAGE EDITOR OPTIONS SPECIFIED XREF,LIST,OVLY,LET
 DEFAULT OPTION(S) USED - SIZE=(196608,65536)

```
IEW0000  INSERT CSA
IEW0000  ENTRY CSA
IEW0000  OVERLAY ALPHA
IEW0000  INSERT CSB
IEW0000  OVERLAY BETA
IEW0000  INSERT CSC
IEW0000  OVERLAY BETA
IEW0000  INSERT CSD
IEW0000  OVERLAY ALPHA
IEW0000  INSERT CSE
IEW0000  OVERLAY GAMMA(REGION)
IEW0000  INSERT CSF
IEW0000  OVERLAY GAMMA
IEW0000  INSERT CSG
IEW0172  2    CSE
IEW0182  4    CSC
```

CROSS REFERENCE TABLE

Root Segment 1:

CONTROL SECTION				ENTRY							
NAME	ORIGIN	LENGTH	SEG. NO.	NAME	LOCATION	NAME	LOCATION	NAME	LOCATION	NAME	LOCATION
\$SEGTAB	00	34	1								
CSA	38	366	1								
ILBODSP0*	3A0	6FB	1								
ILBOSTP0*	A98	35	1								
				ILBOSTP1	AAE						
\$ENTAB	AD0	30	1								
LOCATION	REFERS TO	SYMBOL	IN CONTROL SECTION	SEG. NO.	LOCATION	REFERS TO	SYMBOL	IN CONTROL SECTION	SEG. NO.		
2C0		ILBODSP0	ILBODSP0	1	2C4		ILBOSTP0	ILBOSTP0	1		
2C8		CSG	CSG	7	2CC		CSE	CSE	5		
2D0		CSB	CSB	2	2D4		ILBOSTP1	ILBOSTP0	1		

Segment 2:

CONTROL SECTION				ENTRY							
NAME	ORIGIN	LENGTH	SEG. NO.	NAME	LOCATION	NAME	LOCATION	NAME	LOCATION	NAME	LOCATION
CSB	B00	360	2								
\$ENTAB	E60	18	2								
LOCATION	REFERS TO	SYMBOL	IN CONTROL SECTION	SEG. NO.	LOCATION	REFERS TO	SYMBOL	IN CONTROL SECTION	SEG. NO.		
D54*		ILBODSP0	ILBODSP0	1	D50		ILBOSTP0	ILBOSTP0	1		
D58		CSE	CSE	5	D60		ILBOSTP1	ILBOSTP0	1		
D5C		CSD	CSD	4							

Segment 3:

CONTROL SECTION				ENTRY							
NAME	ORIGIN	LENGTH	SEG. NO.	NAME	LOCATION	NAME	LOCATION	NAME	LOCATION	NAME	LOCATION
CSC	E78	336	3								
LOCATION	REFERS TO	SYMBOL	IN CONTROL SECTION	SEG. NO.	LOCATION	REFERS TO	SYMBOL	IN CONTROL SECTION	SEG. NO.		
10CC		ILBODSP0	ILBODSP0	1	10C8		ILBOSTP0	ILBOSTP0	1		
10D0		ILBOSTP1	ILBOSTP0	1							

Segment 4:

CONTROL SECTION				ENTRY							
NAME	ORIGIN	LENGTH	SEG. NO.	NAME	LOCATION	NAME	LOCATION	NAME	LOCATION	NAME	LOCATION
CSD	E78	362	4								
LOCATION	REFERS TO	SYMBOL	IN CONTROL SECTION	SEG. NO.	LOCATION	REFERS TO	SYMBOL	IN CONTROL SECTION	SEG. NO.		
10CC		ILBODSP0	ILBODSP0	1	10C8		ILBOSTP0	ILBOSTP0	1		
10D4		ILBOSTP1	ILBOSTP0	1	10D0		CSC	CSC	3		

Figure 45 (Part 1 of 2). Linkage Editor Output for Sample Program REGNOVLY

CROSS REFERENCE TABLE

Segment 5:

CONTROL SECTION				ENTRY							
NAME	ORIGIN	LENGTH	SEG. NO.	NAME	LOCATION	NAME	LOCATION	NAME	LOCATION	NAME	LOCATION
CSE	B00	336	5								
LOCATION	REFERS TO SYMBOL	IN CONTROL SECTION	SEG. NO.	LOCATION	REFERS TO SYMBOL	IN CONTROL SECTION	SEG. NO.				
D54	ILBODSP0	ILBODSP0	1	D50	ILBOSTP0	ILBOSTP0	1				
D58	ILBOSTP1	ILBOSTP0	1								

Segment 6:

CONTROL SECTION				ENTRY							
NAME	ORIGIN	LENGTH	SEG. NO.	NAME	LOCATION	NAME	LOCATION	NAME	LOCATION	NAME	LOCATION
CSF	11E0	2FA	6								
LOCATION	REFERS TO SYMBOL	IN CONTROL SECTION	SEG. NO.	LOCATION	REFERS TO SYMBOL	IN CONTROL SECTION	SEG. NO.				
1430	ILBOSTP0	ILBOSTP0	1	1434	ILBOSTP1	ILBOSTP0	1				

Segment 7:

CONTROL SECTION				ENTRY							
NAME	ORIGIN	LENGTH	SEG. NO.	NAME	LOCATION	NAME	LOCATION	NAME	LOCATION	NAME	LOCATION
CSG	11E0	336	7								
LOCATION	REFERS TO SYMBOL	IN CONTROL SECTION	SEG. NO.	LOCATION	REFERS TO SYMBOL	IN CONTROL SECTION	SEG. NO.				
1434	ILBODSP0	ILBODSP0	1	1430	ILBOSTP0	ILBOSTP0	1				
1438	ILBOSTP1	ILBOSTP0	1								
ENTRY ADDRESS	38										
TOTAL LENGTH	1518										
*****	DOES NOT EXIST BUT HAS BEEN ADDED TO DATA SET										
AUTHORIZATION CODE IS	0.										

DIAGNOSTIC MESSAGE DIRECTORY

IEW0172 ERROR - EXCLUSIVE CALL FROM SEGMENT NUMBER PRINTED TO SYMBOL PRINTED.
 IEW0182 ERROR - INVALID EXCLUSIVE CALL FROM SEGMENT NUMBER PRINTED TO SYMBOL PRINTED.

Figure 45 (Part 2 of 2). Linkage Editor Output for Sample Program REGNOVLY

Because the LIST option was specified, the control statement listing is produced. Each control statement is preceded by a special message number, IEW0000.

The control statement listing is followed by two diagnostic message numbers (IEW0172 and IEW0182). The explanation of the messages and the information following each message are given at the end of the output in the diagnostic message directory.

The output for each segment contains a module map and a cross-reference table. The segments are listed as they appear in the overlay structure, top to bottom, left to right, and region by region. (Note that this is also the sequence in which the OVERLAY and INSERT statements must be given.)

Within each segment, a module map lists the control sections in ascending sequence according to their assigned origin. The origin, length, and segment number are listed for each control section, along with any entry names and the location at which each entry name is defined. For example, the root segment has five control sections: \$SEGTAB, which is always the first control section in the root segment; CSA, which is from the object module input; ILBODSP0 and ILBOSTP0, which are from the automatic call library (indicated by an asterisk) and were not repositioned; and \$ENTAB, which, when present, is always the last control section in any segment (as also in segment 2). One entry name is defined, ILBOSTP1 at location D58 in control section ILBOSTP0.

The cross-reference table for each segment contains all the address constants that refer to symbols defined in other control sections. The location of the address constant is followed by the symbol referred to, the control section in

which the symbol is defined, and the segment in which the control section is located. For example, in the root segment, an address constant at location 11E0 refers to symbol CSG, which is defined in control section CSG in segment 7. Although the region is not given, the overlay tree in Figure 43 on page 128 shows that segment 7 is in region 2.

At the end of the output for all the segments are the entry address and total length. The entry address is 38, which is the origin of CSA, the specified entry point. The total length given refers to main storage used, not device storage. The length given, therefore, is that of the longest path. The longest path is that formed by the root segment and segments 2, 4, and 7; the length given is 1518.

However, if the given lengths of the control sections in each segment are added, the result is 14D3. The discrepancy exists because the given lengths do not include the padding bytes necessary to make control sections begin on a doubleword address (multiple of 8). For example, in the root segment, the length of \$SEGTAB is 34; however, the origin of CSA which follows \$SEGTAB is 38 (decimal 56). Four additional bytes are needed so that the origin of CSA is a multiple of 8.

The disposition message indicates that the load module GO has been added to the output module library. The library did not contain any other module by that name. The four asterisks identify the message.

The last item in the output for this sample program is the diagnostic message directory. The directory contains the text for the message numbers listed after the control statement listing. The directory must be correlated to the information following the number to interpret the message.

For example, message IEW0172 is an error message that indicates that an exclusive call was made from the segment number printed (2) following the message number to the symbol printed (CSE). The output for segment 2 indicates that this call is at location D58 in control section CSB, and the symbol is defined in control section CSE in segment 5. This is the valid exclusive call from CSB to CSE described earlier. (If XCAL were specified, a warning message would be issued instead of an error message.)

If an invalid exclusive call is detected, message IEW0182 appears as shown. This is also an error message; it indicates that an invalid exclusive call was made from segment 4 to symbol CSC. This call is at location E78 in control section CSD, and the symbol is defined in control section CSC in segment 3. This is the invalid exclusive call from CSD to CSC, also described earlier.

SAMPLE PROGRAM PARTDS

Sample program PARTDS illustrates that linkage editor control statements can be placed in a separate data set and then used as input. For convenience, the control statements are those for sample program REGNOVLY, described previously. These control statements are placed in a partitioned data set. When the member that contains the control statements is referenced, the linkage editor uses the control statements to produce the overlay structure shown in Figure 43 on page 128.

Figure 46 on page 133 shows the input statements for the IEBUPDTE utility program used to place the control statements in a partitioned data set.

```

//PARTDS    JOB      (accounting information)
//CTLG      EXEC     PGM=IEBUPDTE,PARM=(NEW)
//SYSUT2    DD       DSNAME=OVLYLIB,UNIT=2314,VOL=SER=DA028,DISP=(NEW,KEEP),
//          //       SPACE=(TRK,(10,5,2)),DCB=(LRECL=80,BLKSIZE=80,RECFM=F)
//SYSPRINT  DD       SYSOUT=A
//SYSIN     DD       *
./          ADD      NAME=OVLY,LEVEL=00,SOURCE=00,LIST=ALL
./          NUMBER   NEW1=10,INCR=5
  INSERT CSA
  ENTRY CSA
  OVERLAY ALPHA
  INSERT CSB
  OVERLAY BETA
  INSERT CSC
  OVERLAY BETA
  INSERT CSD
  OVERLAY ALPHA
  INSERT CSE
  OVERLAY GAMMA(REGION)
  INSERT CSF
  OVERLAY GAMMA
  INSERT CSG
  /          ENDUP
/*

```

Figure 46. Input Statements for IEBUPDTE Utility Program

The source programs for all the control sections were compiled in previous job steps. All the object modules were placed in the same data set, which was passed to the linkage editor job step. The input modules are those used for sample program REGNOVLY.

Job Control Language

The job control language for the overlay program job step of this sample program is shown below.

```

//LKED      EXEC     PGM=HEWL,PARM='XREF,LIST,OVLY,LET'
//SYSUT1    DD       DSNAME=&&UT1,UNIT=SYSDA,SPACE=(TRK,
//          //       (100,10))
//OVLYCDS   DD       DSNAME=OVLYLIB,UNIT=SYSDA,
//          //       VOL=SER=SCRATCH,DISP=OLD
//SYSLIB    DD       DSNAME=SYS1.COBLIB,DISP=SHR
//SYSLMOD   DD       DSNAME=&&OVLYJB(GO),UNIT=SYSDA,
//          //       DISP=(NEW,PASS),SPACE=(TRK,(100,10,1))
//SYSPRINT  DD       SYSOUT=A
//SYSLIN    DD       DSNAME=&&OBJMOD,DISP=(OLD,DELETE)
//          DD       *

```

(Linkage Editor Control Statements)

/*

Statement	Explanation
EXEC	Causes the execution of the linkage editor. The PARM field options request a cross-reference table and a module map (XREF), and a control statement listing (LIST) to be produced on the diagnostic output data set. The output load module is to be assigned the overlay attribute (OVLY), and is to be marked executable despite severity 2 errors (LET).
SYSUT1	Defines a temporary direct access data set to be used as the intermediate data set.
OVLYCDS	Defines a permanent data set to be used later as additional input; this is the partitioned data set which was created by IEBUPDTE and contains the control statements for structuring the overlay program.
SYSLIB	Defines the automatic call library (SYS1.COBLIB) to be used to resolve external references. All control sections from this library are placed in the root segment; they remain there unless they are repositioned.
SYSLMOD	Defines a temporary data set to be used as the output module library; the load module is to be assigned the member name GO, and is passed to a subsequent step for execution.
SYSPRINT	Defines the diagnostic output data set, which is assigned to output class A.
SYSLIN	Defines the primary input data set, &&OBJMOD, which contains the object modules for the overlay structure. This data set is temporary and was passed from a previous job step; it is to be deleted at the end of this job. This statement also concatenates the input stream to the primary input data set. The input stream contains linkage editor control statements that must be delimited by a /* statement.

Figure 47. Job Control Statements for PARTDS

Linkage Editor Control Statements

The input stream contains an INCLUDE statement, as follows:

```
INCLUDE OVLYCDS(OVLY)
```

This statement causes the control statements to be read from the partitioned data set described on the OVLYCDS DD statement. The member name of the statements is OVLY, the same name used in the ADD statement for the utility program.

Linkage Editor Output

The output of this sample program is identical to the output from the REGNOVLY sample program, with one exception. The list of control statements begins with the statement

```
IEW0000    INCLUDE  OVLYCDS(OVLY)
```

This statement is followed by a list of the control statements read from the additional input data set specified in this INCLUDE statement. The rest of the output is identical to that shown in Figure 45 on page 130.

APPENDIX B. LINKAGE EDITOR REQUIREMENTS AND CAPACITIES

This appendix describes the record-processing capacities of the linkage editor, the types of devices that can be used for the intermediate data set (SYSUT1), and the amount of virtual storage the linkage editor requires.

Capacities

The minimum storage requirement and processing capacities of the linkage editor program are described in Figure 44 on page 129. To increase the capacity for processing external symbol dictionary records, intermediate text records, relocation dictionary records, and identification records, increase value1 and/or value2 of the SIZE option. Output text record length can be increased by increasing the SIZE option values, but in no case may the record length ever exceed the lowest track length for the device or 18K bytes. The number of overlay segments and regions that can be processed is not affected by increasing the storage available.

Function	Capacity (Bytes)
Virtual storage allocated (in bytes)	96K
Maximum number of entries in composite external symbol dictionary (CESD)	558
Maximum number of intermediate text records	372
Maximum number of relocation dictionary (RLD) records	192
Maximum number of segments per program	255
Maximum number of overlay regions per program	4
Maximum blocking factor for input object modules object modules (number of 80-column card images per physical record)	5
Maximum blocking factor for SYSPRINT output (number of 121-character logical records per physical record)	5
Output text record length (in bytes), for the devices supported by this system: 2305-2 Fixed Head Storage, 3330 Disk Storage, 3330-11 Disk Storage, 3340 DASD, 3344 DASD, 3350 DASD, 3375 DASD, 3380 DASD ² .	3072 ¹

Figure 48. Linkage Editor Capacities for Minimal SIZE Values (96K bytes, 6K bytes)

Notes to Figure 48:

¹ The maximum output text record length is achieved when value2 of the SIZE parameter is at least twice the record length size. For example, on a 3330, 12288 byte records are written when value2 is at least 24576.

² 3380 Models A04, AA4, B04, AD4, BD4, AE4, and BE4.

For the composite external symbol dictionary, the number of entries permitted can be computed by subtracting, from the maximum number given in Figure 48 on page 136, one entry for each of the following:

- A data definition name (ddname) specified in LIBRARY statements
- A data definition name (ddname) specified in INCLUDE statements
- An ALIAS statement
- A symbol in REPLACE or CHANGE statements that are in the largest group of such statements preceding a single object module in the input to the linkage editor
- The segment table (SEGTAB) in an overlay program
- An entry table (ENTAB) in an overlay program

To compute the number of intermediate text records that will be produced during processing of either program, add one record for each group of x bytes within each control section, where x is the record size for the intermediate data set. The minimum value for x is 1024; a maximum is chosen depending on the amount of storage available to the linkage editor and the devices allocated for the intermediate and output data sets.

The number of intermediate text records that can be handled by a linkage editor program is less than the maximums given in Figure 48 on page 136 if the text of one or more control sections is not in sequence by address in the input to the linkage editor.

The total length of the data fields of the CSECT identification records associated with a load module cannot exceed 32K (32768) bytes. To determine the number of bytes of identification data contained in a particular load module, use the following formula:

$$\text{SIZE} = 269 + 16A + 31B + 2C + I(n + 6)$$

where:

- A = the number of compilations or assemblies by a processor supporting CSECT identification that produced the object code for the module.
- B = the number of preprocessor compiler compilations by a processor supporting CSECT identification that produced the object code for the module.
- C = the number of control sections in the module with END statements that contain identification data.
- I = the number of control sections in the module that contain user-supplied data supplied during link-editing by the optional IDENTIFY control statement.
- n = the average number of characters in the data specified by IDENTIFY control statements.

Notes:

1. The size computed by the formula includes space for recording up to 19 HMASPZAP modifications. When 75% of this space has been used, a new 251-byte record is created the next time the module is reprocessed by the linkage editor.
2. To determine the approximate number of records involved, divide the computed size of the identification data by 256.

EXAMPLE: A module contains 100 control sections produced by 20 unique compilations. Each control section is identified during link-editing by 8 characters of user data specified by the IDENTIFY control statement. The size of the identification data is computed as follows:

A = 20
I = 100
n = 8

$269 + 320 + 1400 = 1989$ bytes

If the optional user data specified on the IDENTIFY control statements is omitted, the size can be reduced considerably, as computed below:

$269 + 320 = 589$ bytes

The maximum number of downward calls made from a segment to other segments lower in its path can never exceed 340. To compute the maximum number of downward calls allowed, subtract 12 from the SYSLMOD record size and then divide the difference by 12. Examples of maximum downward calls are 84 for a SYSLMOD record size of 1024 bytes and 340 for a SYSLMOD record size of 6144 bytes.

Intermediate Data Set

The intermediate data set (SYSUT1) is used by the linkage editor to hold intermediate data records during processing. The linkage editor places intermediate data in this data set when storage allocated for input data or certain forms of out-of-sequence text is exhausted.

The following direct access devices, if supported by the system, can be used for this data set:

IBM 2305-1	Fixed Head Storage Facility
IBM 2305-2	Fixed Head Storage
IBM 2314	Storage Control
IBM 2319	Disk Storage
IBM 3330	Disk Storage
IBM 3330-11	Disk Storage
IBM 3340	Direct Access Storage
IBM 3344	Direct Access Storage
IBM 3350	Direct Access Storage
IBM 3375	Direct Access Storage
IBM 3380 ¹	Direct Access Storage

¹ 3380 models A04, AA4, B04, AD4, BD4, AE4, and BE4.

APPENDIX C. DESIGNING AND SPECIFYING OVERLAY PROGRAMS

Ordinarily, when a load module produced by the linkage editor is executed, all the control sections of the module remain in virtual storage throughout execution. The length of the load module is, therefore, the sum of the lengths of all the control sections. When storage space is not at a premium, this is the most efficient way to execute a program. However, if a program approaches the limits of the virtual storage available, the programmer should consider using the overlay facilities of the linkage editor.

In most cases, all that is needed to convert an ordinary program to an overlay program is the addition of control statements to structure the module. The programmer chooses the overlayable portions of the program, and the system arranges to load the required portions when needed during execution of the program.

When the linkage editor overlay facility is requested, the load module is structured so that, at execution time, certain control sections are loaded only when referenced. When a reference is made from an executing control section to another, the system determines whether or not the code required is already in virtual storage. If it is not, the code is loaded dynamically and may overlay an unneeded part of the module already in storage.

The rest of this chapter is divided into three sections that describe the design, specification, and special considerations for overlay programs.

DESIGN OF AN OVERLAY PROGRAM

The way in which an overlay module is structured depends on the relationships among the control sections within the module. Two control sections that do not have to be in storage at the same time can overlay each other. Such control sections are independent; that is, they do not reference each other either directly or indirectly. Independent control sections can be assigned the same load addresses and are loaded only when referenced. For example, control sections that handle error conditions or unusual data may be used infrequently, and need not be occupying storage unless in use.

Control sections are grouped into segments. A segment is the smallest functional unit (one or more control sections) that can be loaded as one logical entity during execution. The control sections required all the time are grouped into a special segment called the root segment. This segment remains in storage throughout execution of an overlay program.

When a particular segment is to be executed, any segments between it and the root segment must also be in storage. This is a path. A reference from one segment to another segment lower in a path is a downward reference (see "Control Section Dependency" on page 140). That is, the segment contains a reference to another segment farther from the root segment. Conversely, a reference from one segment to another segment higher in a path (closer to the root segment) is an upward reference.

Therefore, a downward reference may cause overlay, because the necessary segment may not yet be in virtual storage. An upward reference will not cause overlay, because all segments between a segment and the root segment must be present in storage.

Sometimes several paths need the same control sections. This problem may be solved by placing the control sections in another region. In an overlay structure, a region is a contiguous area of virtual storage within which segments can be loaded independently of paths in other regions. An overlay program can be designed in single or multiple regions.

SINGLE REGION OVERLAY PROGRAM

To design an overlay structure, the programmer should select those control sections that will receive control at the beginning of execution, plus those that should always remain in storage; these control sections form the root segment. The rest of the structure is developed by determining the dependencies of the remaining control sections and how they can use the same virtual storage locations at different times during execution.

Besides control section dependency, other topics discussed in this section are segment dependency, the length of the overlay program, segment origin, communication between segments, and overlay processing.

Control Section Dependency

Control section dependency is determined by the requirements of a control section for a given routine in another control section. A control section is dependent upon any control section from which it receives control, or which processes its data. For example, if control section C receives control from control section B, then C is dependent upon B. That is, both control sections must be in storage before execution can continue beyond a given point in the program.

Assume that a program contains seven control sections, CSA through CSG, and exceeds the amount of storage available for its execution. Before the program is rewritten, it is examined to see whether or not it could be placed into an overlay structure. Figure 49 on page 141 shows the groups of dependent control sections in the program (the arrows indicate dependencies).

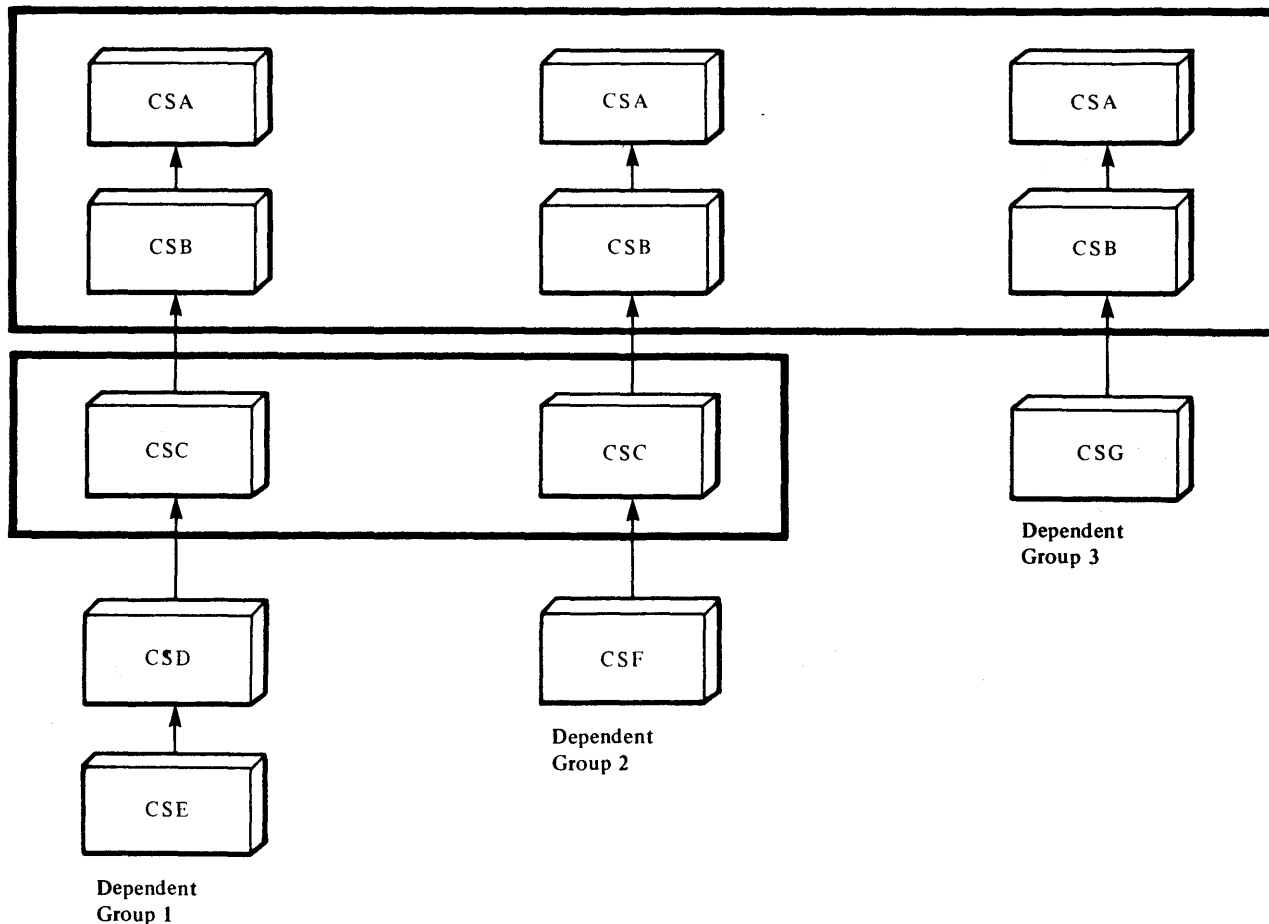


Figure 49. Control Section Dependencies

Each dependent group is also a path. That is, if control section CSG is to be executed, CSB and CSA must also be in storage. Because CSA and CSB are in each path, they must be in the root segment. Control section CSC is in two groups, and therefore is a common segment in two different paths.

A better way to show the relationship between segments is with a tree structure. A tree is the graphic representation that shows how segments can use virtual storage at different times. It does not imply the order of execution, although the root segment is the first to receive control. Figure 50 on page 142 shows the tree structure for the dependent groups shown in Figure 49. The structure is contained in one region, and has five segments.

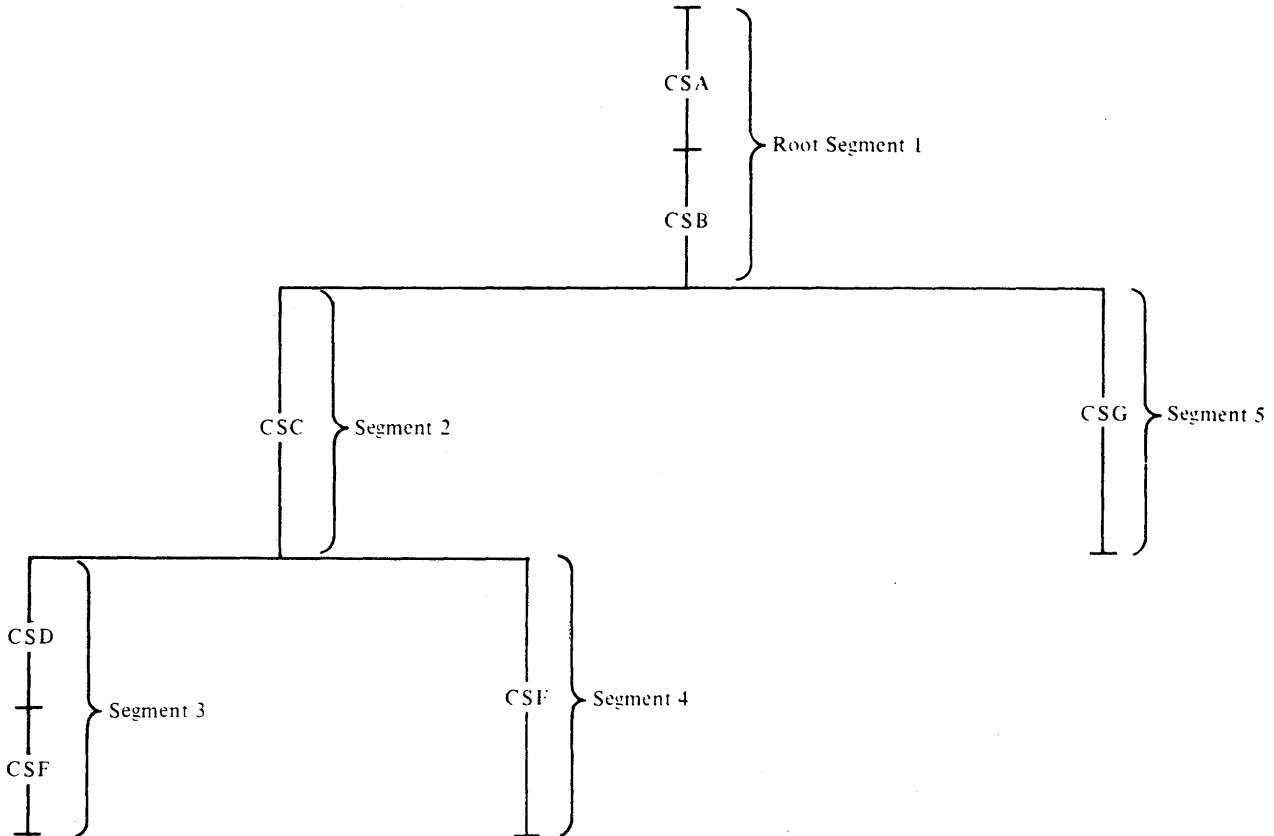


Figure 50. Single-Region Overlay Tree Structure

Segment Dependency

When a segment is in virtual storage, all segments in its path are also in virtual storage. Each time a segment is loaded, all segments in its path are loaded if they are not already in virtual storage. In Figure 50, when segment 3 is in virtual storage, segments 1 and 2 are also in virtual storage. However, if segment 2 is in storage, this does not imply that segment 3 or 4 is in virtual storage, because neither segment is in the path of segment 2.

The position of the segments in an overlay tree structure does not imply the sequence in which the segments are executed. A segment can be loaded and overlaid as many times as required by the logic of the program. However, a segment will not be overlaid by itself. If a segment is modified during execution, that modification remains only until the segment is overlaid.

Length of an Overlay Program

For purposes of illustration, assume that the control sections in the sample program have the following lengths:

Control Section	Length (in bytes)
CSA	3000
CSB	2000
CSC	6000
CSD	4000
CSE	3000
CSF	6000
CSG	8000

If the program were not in overlay, it would require 32000 bytes of virtual storage. In overlay, however, the program requires the amount of storage needed for the longest path. In this structure, the longest path is formed by segments 1, 2, and 3, since, when they are all in storage, they require 18000 bytes, as shown in Figure 51.

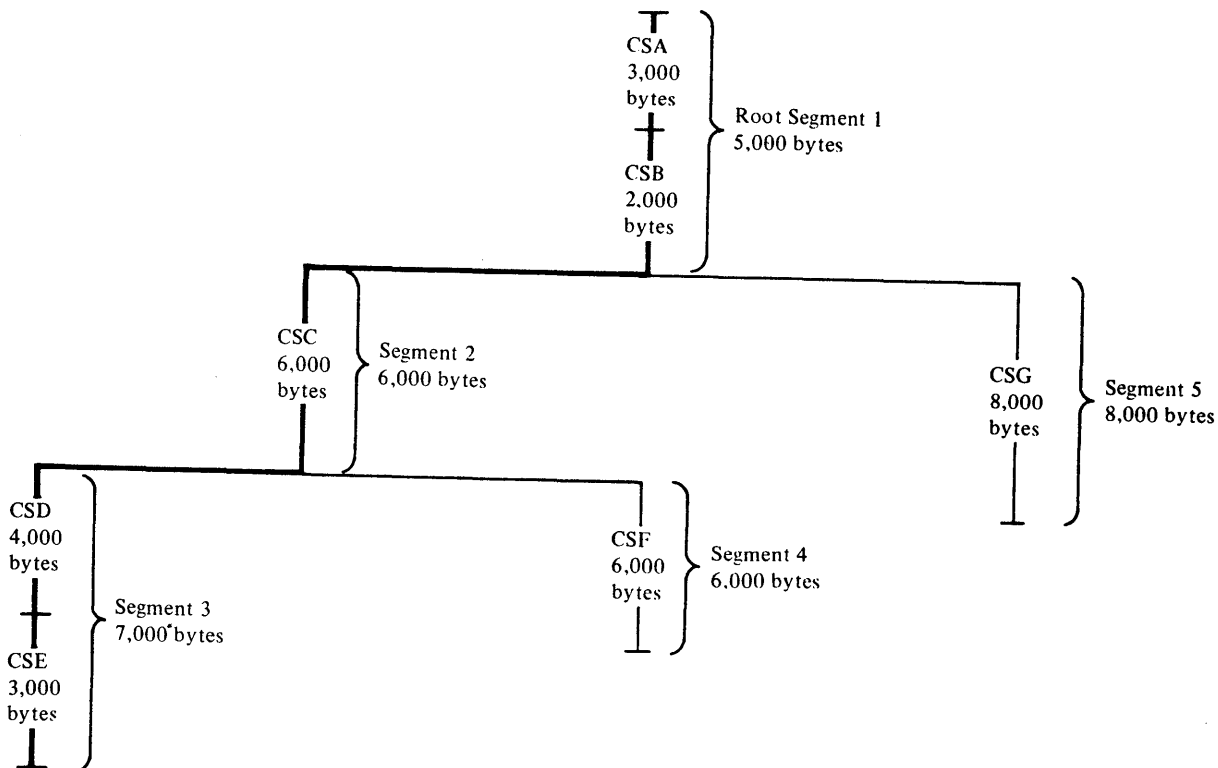


Figure 51. Length of an Overlay Module

Note: The length of the longest path is not the minimum requirement for an overlay program; when a program is in overlay, certain tables are used, and their storage requirements must also be considered. The storage required by these tables is given in the section "Special Considerations" on page 157.

Segment Origin

The linkage editor assigns the relocatable origin of the root segment (the origin of the program) at 0. The relative origin of each segment is determined by 0 plus the length of all segments in the path. For example, the origin of segments 3 and 4 is equal to 0 plus 6000 (the length of segment 2) plus 5000 (the length of the root segment), or 11000. The origins of all the segments are as follows:

Segment	Origin
1	0
2	5000
3	11000
4	11000
5	5000

The segment origin is also called the load point, because it is the relative location at which the segment is loaded.

Figure 52 shows the segment origin for each segment and the way storage is used by the sample program. In the illustration, the vertical bars indicate segment origin; any two segments with the same origin may use the same storage area. Figure 52 also shows that the longest path is that of segments 1, 2, and 3.

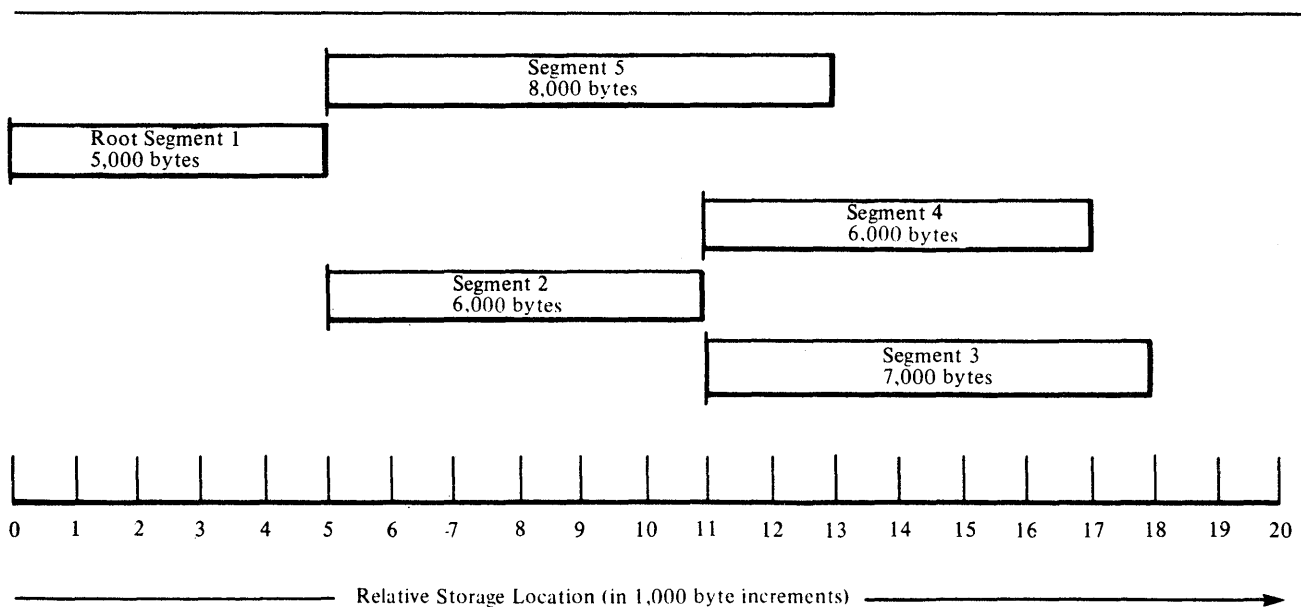


Figure 52. Segment Origin and Use of Storage

Communication between Segments

Segments that can be in virtual storage simultaneously are considered to be inclusive. Segments in the same region but not in the same path are considered to be exclusive; they cannot be in virtual storage simultaneously. Figure 53 shows the inclusive and exclusive segments in the sample program.




Figure 53. Inclusive and Exclusive Segments

Segments upon which two or more exclusive segments are dependent are called common segments. A segment common to two other segments is part of the path of each segment. Figure 53, segment 2 is common to segments 3 and 4, but not to segment 5.

An inclusive reference is a reference between inclusive segments; that is, a reference from a segment in storage to an external symbol in a segment that will not cause overlay of the calling segment. An exclusive reference is a reference between exclusive segments; that is, a reference from a segment in storage to an external symbol in a segment that will cause overlay of the calling segment.

Figure 54 on page 146 shows the difference between an inclusive reference and an exclusive reference; the arrows indicate references between segments.

INCLUSIVE REFERENCES: Wherever possible, inclusive references should be used instead of exclusive references. Inclusive references between segments are always valid and do not require special options. When inclusive references are used, there is also less chance for error in structuring the overlay program correctly.

EXCLUSIVE REFERENCES: An exclusive reference is made when the external reference in the requesting segment is to a symbol defined in a segment not in the path of the requesting segment. Exclusive references are either valid or invalid.

An exclusive reference is valid only if there is also an inclusive reference to the requested control section in a segment common to both the segment to be loaded and the segment to be overlaid. The same symbol must be used in both the common segment and the exclusive reference. In Figure 54, a reference from segment B to segment A is valid, because there is an inclusive reference from the common segment to segment A. (An entry table in the common segment contains the address of segment A; the overlay does not destroy this table.)




Figure 54. Inclusive and Exclusive References

In the same illustration, a reference from segment A to segment B is invalid, because there is no reference from the common segment to segment B. A reference from segment A to segment B can be made valid by including, in the common segment, an external reference to the symbol used in the exclusive reference to segment B.

Another way to eliminate exclusive references is to arrange the program so that the references that will cause overlay are made in a higher segment. For example, the programmer could eliminate the exclusive reference shown in Figure 54 by writing a new module to be placed in the common segment; the new module's only function would be to reference segment B. The code in segment A could then be changed to refer to the new module instead of to segment B. Control then would pass from segment A to the common segment, where the overlay of segment A by segment B would be initiated.

If either valid or invalid exclusive references appear in the program, the linkage editor considers them errors unless one of the special options is used. These options are described later in this section (see "Special Considerations" on page 157).

Notes:

1. During the execution of a program written in a higher level language such as FORTRAN, COBOL, or PL/I, an exclusive call results in abnormal termination of the program if the requested segment attempts to return control directly to the invoking segment that has been overlaid.
2. If a program written in COBOL includes a segment that contains a reference to a COBOL class test or TRANSFORM

table, the segment containing the table must be either (1) in the root segment or (2) a segment that is higher in the same path than the segment containing the reference to the table.

Overlay Process

The overlay process is initiated during execution of a program only if a control section in virtual storage references a control section not in storage. The control program determines the segment that the referenced control section is in and, if necessary, loads the segment. When a segment is loaded, it overlays any segment in storage with the same relative origin. Any segments in storage that are lower in the path of the overlaid segment may also be overlaid. An exclusive reference can also cause segments higher in the path to be overlaid. If a control section in storage references a control section in another segment already in storage, no overlay occurs.

The portion of the control program that determines when overlay is to occur is the overlay supervisor, which uses special tables to determine when overlay is necessary. These tables are generated by the linkage editor, and are part of the output load module. The special tables are the segment table and the entry table(s). Figure 55 shows the location of the segment and entry tables in the sample program.

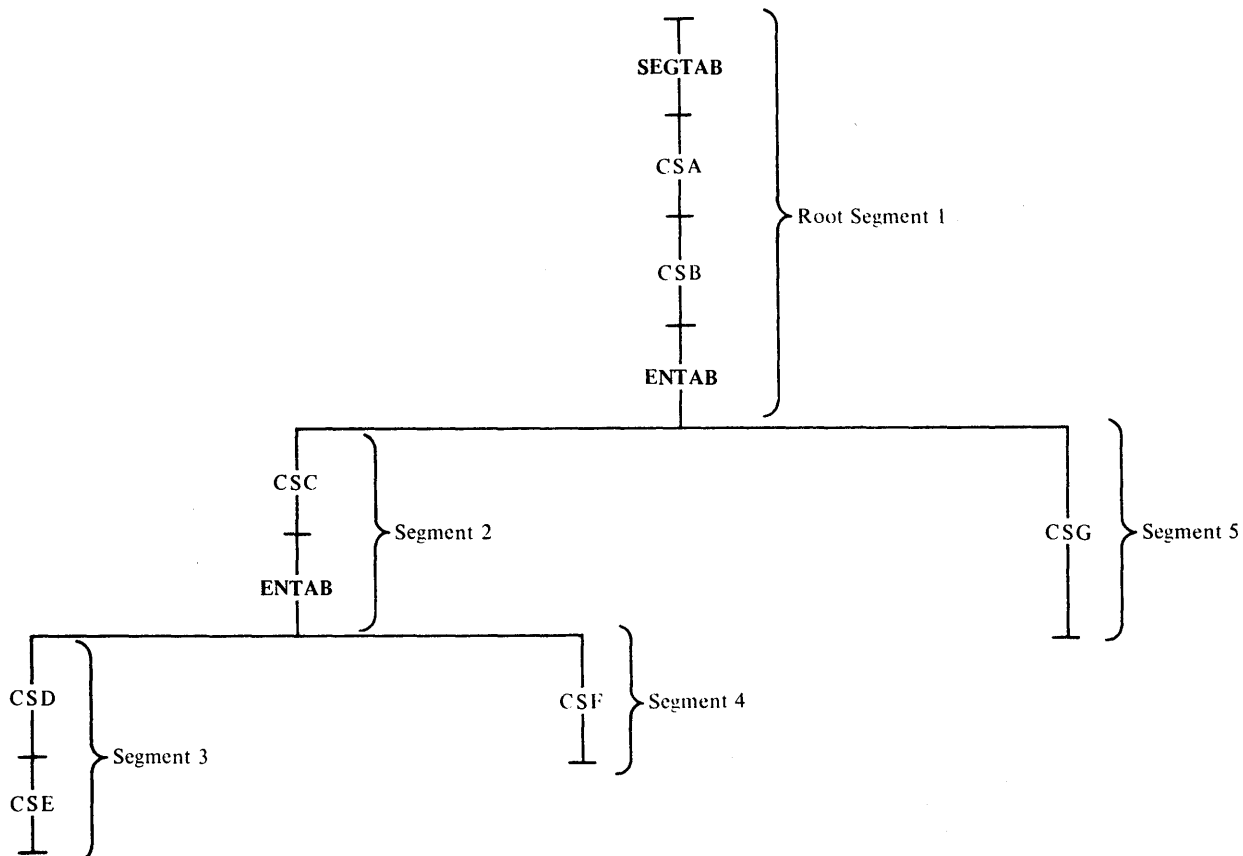


Figure 55. Location of Segment and Entry Tables in an Overlay Module

Because the tables are present in every overlay module, their size must be considered when planning the use of virtual

storage. The storage requirements for the tables are given in "Special Considerations" on page 157. A more detailed discussion of the segment and entry tables follows.

SEGMENT TABLE: Each overlay program contains one segment table (SEGTAB); this table is the first control section in the root segment. The segment table contains information about the relationship of the segments and regions in the program. During execution, the table also indicates which segments are either in storage or being loaded, and other control information.

ENTRY TABLE: Each segment that is not the last segment in a path may contain one entry table (ENTAB); this table, when present, is the last control section in a segment.

When overlay will be required, an entry in the table is created for a symbol to which control is to be passed, provided (1) the symbol is used as an external reference in the requesting segment, and (2) the symbol is defined in another segment either lower in the path of the requesting segment, or in another region. An ENTAB entry is not created for any symbol already present in an entry table closer to the root segment (higher in the path), or for a symbol defined higher in the path. (A reference to a symbol higher in the path does not have to go through the control program because no overlay is required.)

If an external reference and the symbol to which it refers are in segments not in the same path but in the same region, an exclusive reference was made. If the exclusive reference is valid, an ENTAB entry for the symbol is present in the common segment. Because the common segment is higher in the path of the requesting segment, no ENTAB entry is created in the requesting segment. When the reference is executed, control passes through the ENTAB entry in the common segment. That is, a branch to the location in the ENTAB entry causes the overlay supervisor to be called to load the needed segment or segments.

If the exclusive reference is invalid, no ENTAB entry is present in the common segment. If the LET option is specified, an invalid exclusive reference causes unpredictable results when the program is executed. Because no ENTAB entry exists, control is passed directly to the relative address specified in the reference, even though the requested segment may not be in virtual storage.

MULTIPLE REGION OVERLAY PROGRAM

If a control section is used by several segments, it is usually desirable to place that control section in the root segment. However, the root segment can get so large that the benefits of overlay are lost. If some of the control sections in the root segment could overlay each other (except for the requirement that all segments in a path must be in storage at the same time), the job may be a candidate for multiple region structure. Multiple region structures can also be used to increase segment loading efficiency: Processing can continue in one region while the next path to be executed is being loaded into another region.

With multiple regions, a segment has access to segments that are not in its path. Within each region, the rules for single region overlay programs apply, but the regions are independent of each other. A maximum of four regions can be used.

Figure 56 on page 149 shows the relationship between the control sections in the sample program and two new control sections, CSH and CSI. The two new control sections are each used by two other control sections in different paths. Placing CSH and CSI in the root segment makes the segment larger than necessary, because CSH and CSI can overlay each other. The two control sections should not be duplicated in two paths, because

the linkage editor automatically deletes the second pair and an invalid exclusive reference may then result.




Figure 56. Control Sections Used by Several Paths

If, however, the two control sections are placed in another region, they can be in virtual storage when needed, regardless of the path being executed in the first region. Figure 57 on page 150 shows all the control sections in a two-region structure. Either path in region 2 can be in virtual storage regardless of the path being executed in region 1; segments in region 2 can cause segments in region 1 to be loaded without being overlaid themselves.

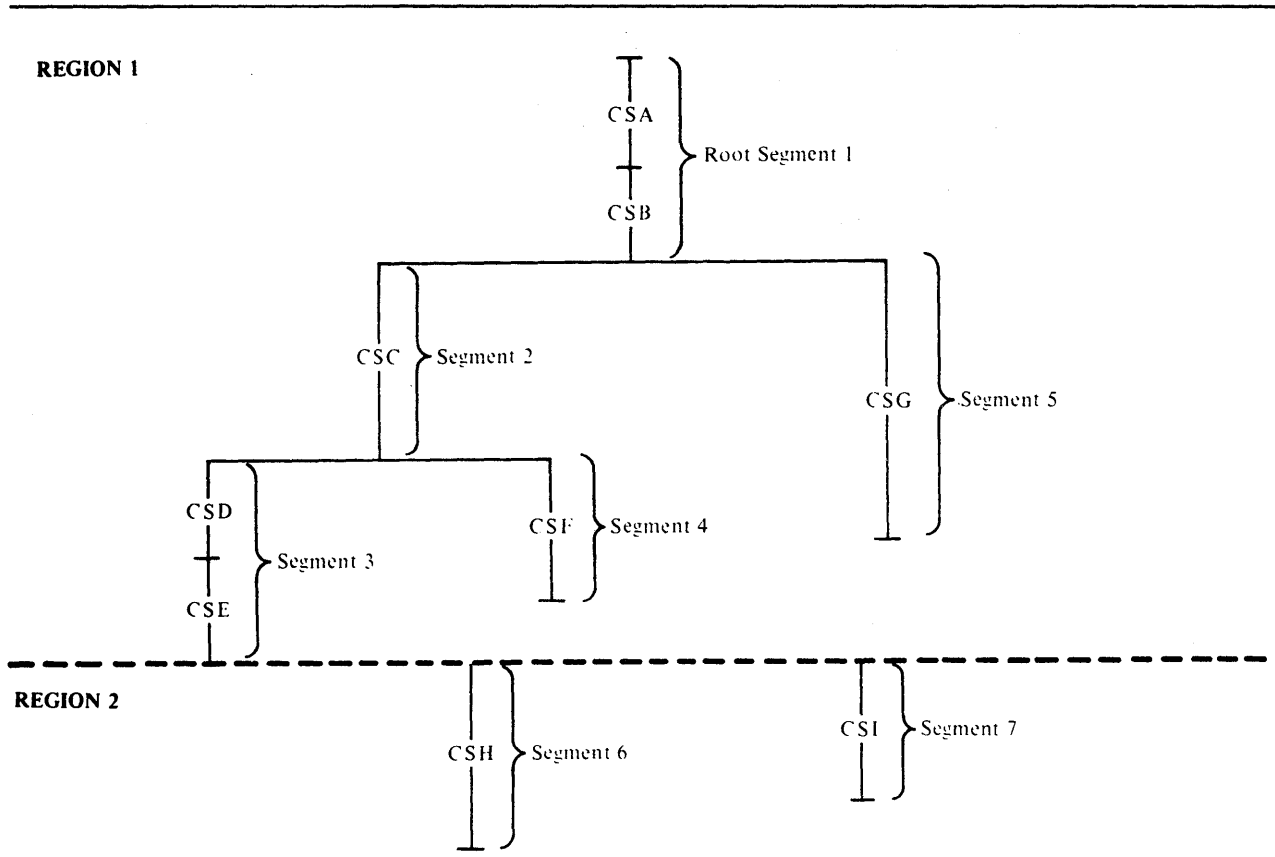


Figure 57. Overlay Tree for Multiple-Region Program

The relative origin of a second region is determined by the length of the longest path in the first region (18000 bytes). Region 2, therefore, begins at 0 plus 18000 bytes. The relative origin of a third region would be determined by the length of the longest path in the first region plus the longest path in the second region.

The virtual storage required for the program is determined by adding the lengths of the longest path in each region. In Figure 57, if CSH is 4000 bytes and CSI is 3000 bytes, the storage required is 22000 bytes, plus the storage required by the special overlay tables.

Care should be exercised when choosing multiple regions. There may be some system degradation caused by the overlay supervisor being unable to optimize segment loading when multiple regions are used.

SPECIFICATION OF AN OVERLAY PROGRAM

Once the programmer has designed an overlay structure, the module must be placed in that structure by indicating to the linkage editor the relative positions of the segments and regions, and the control sections in each segment. Positioning is accomplished as follows:

- Segments are positioned by OVERLAY statements. In addition, the overlay statement provides a means by which to equate each load point with a unique symbolic name. Each OVERLAY statement begins a new segment.

- Regions are also positioned by OVERLAY statements. The programmer specifies the origin of the first segment of the region, followed by the word REGION in parentheses.
- Control sections are positioned in the segment specified by the OVERLAY statement with which they are associated in the input sequence. However, the sequence of the control sections within a segment is not necessarily the order in which the control sections are specified.

The input sequence of control statements and control sections should reflect the sequence of the segments in the overlay structure from top to bottom, left to right, and region by region. This sequence is illustrated in later examples.

In addition, several special options are used with overlay programs. These options are specified on the EXEC statement for the linkage editor job step, and are described at the end of this section.

Note: If a load module in overlay structure is to be reprocessed by the linkage editor, the OVERLAY statements and special options (such as OVLY) must be respecified. If the statements and options are not provided, the output load module will not be in overlay structure.

SEGMENT ORIGIN

The symbolic origin of every segment, other than the root segment, must be specified with an OVERLAY statement. The first time a symbolic origin is specified, a load point is created at the end of the previous segment. That load point is logically assigned a relative address at the doubleword boundary that follows the last byte in the preceding segment. Subsequent use of the same symbolic origin indicates that the next segment is to have its origin at the same load point.

In the sample single-region program, the symbolic origin names ONE and TWO are assigned to the two necessary load points, as shown in Figure 57 on page 150. Segments 2 and 5 are at load point ONE; segments 3 and 4 are at load point TWO.

The following sequence of OVERLAY statements will result in the structure in Figure 58 on page 152 (the control sections in each segment are indicated by name):

```
Control section CSA
Control section CSB
OVERLAY ONE
Control section CSC
OVERLAY TWO
Control section CSD
Control section CSE
OVERLAY TWO
Control section CSF
OVERLAY ONE
Control section CSG
```

Note: The sequence of OVERLAY statements reflects the order of segments in the structure from top to bottom and left to right.

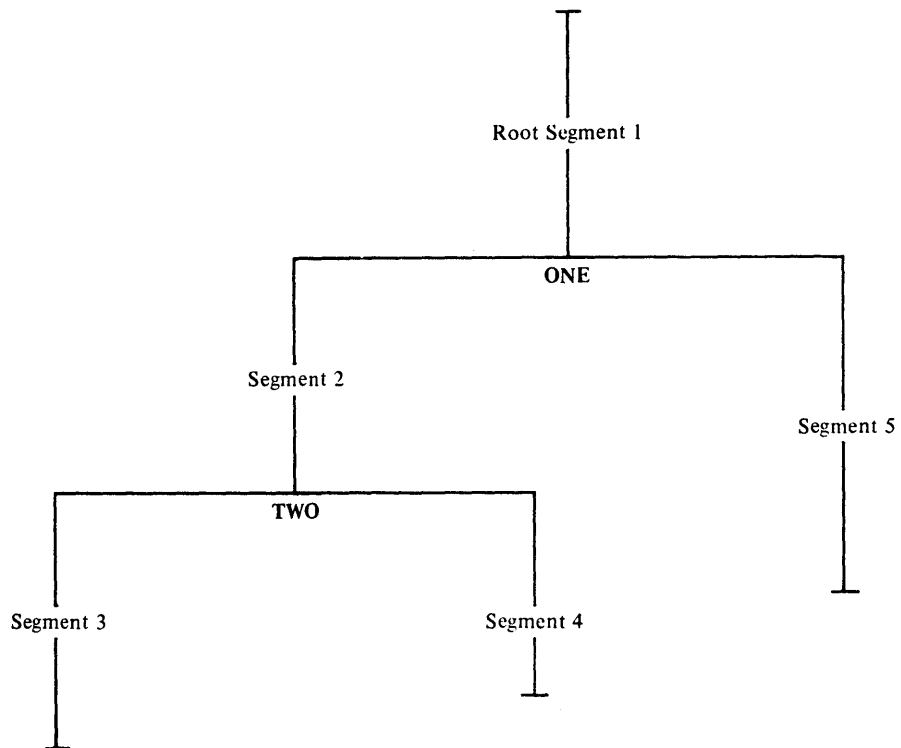


Figure 58. Symbolic Segment Origin in Single-Region Program

REGION ORIGIN

The symbolic origin of every region, other than the first, must be specified with an OVERLAY statement. Once a new region is specified, a segment origin from a previous region should not be specified.

In the sample multiple-region program, the symbolic origin THREE is assigned to region 2, as shown in Figure 59 on page 153. Segments 6 and 7 are at load point THREE.

REGION 1

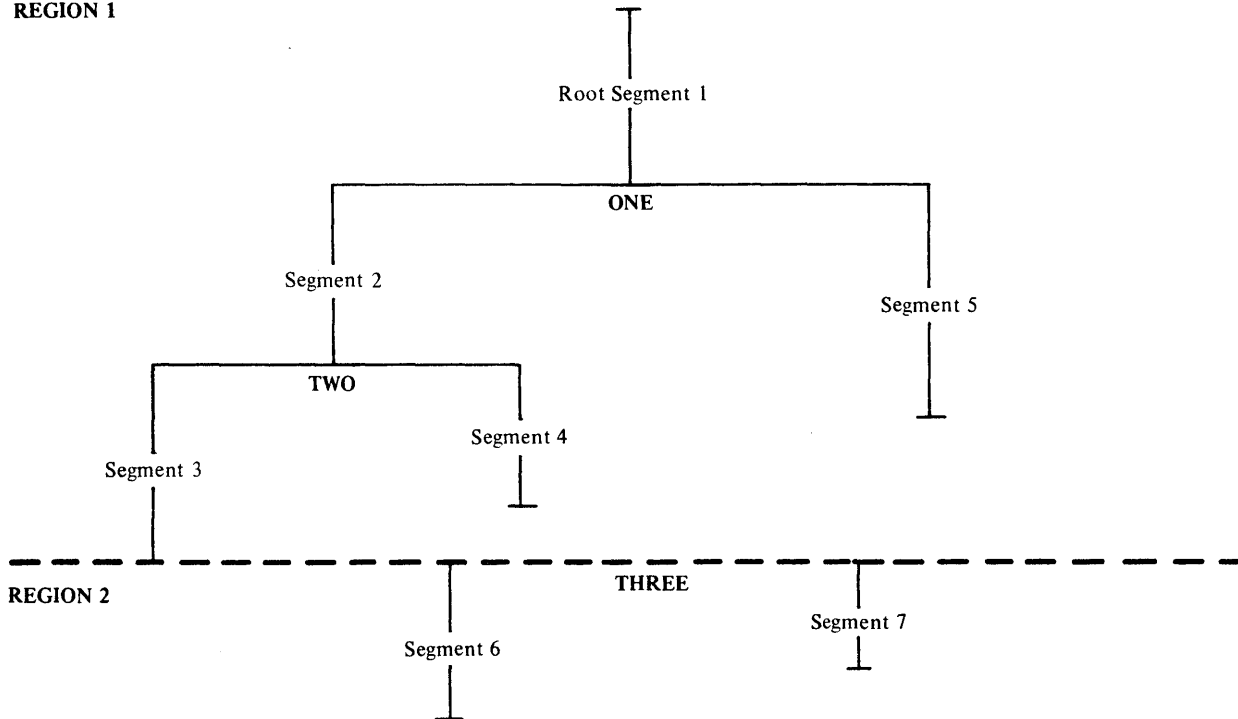


Figure 59. Symbolic Segment and Region Origin in Multiple-Region Program

If the following is added to the sequence for the single-region program, the multiple-region structure will be produced:

```

:
:
OVERLAY THREE(REGION)
Control section CSH
OVERLAY THREE.
Control section CSI
```

POSITIONING CONTROL SECTIONS

After each OVERLAY statement, the control sections for that segment must be specified. The control sections for a segment can be specified in one of three ways:

- By placing the object decks for each segment after the appropriate OVERLAY statement
- By using INCLUDE control statements for the modules containing the control sections for the segment
- By using INSERT control statements to reposition a control section from its position in the input stream to a particular segment

Any control sections that precede the first OVERLAY statement are placed in the root segment; they can be repositioned with an INSERT statement. Control sections from the automatic call library are also placed in the root segment. The INSERT

statement can be used to place these control sections in another specific segment. Common areas in an overlay program are described in "Special Considerations" on page 157.

An example of each of the three methods of positioning control sections follows. Each example results in the structure for the single-region sample program. An example is also given of repositioning control sections from the automatic call library.

Using Object Decks

The primary input data set for this example contains an ENTRY statement and seven object decks, separated by OVERLAY statements:

```
//LKED          EXEC  PGM=HEWL,PARM='OVLY'  
.  
.  
//SYSLIN        DD    *  
ENTRY BEGIN.  
Object deck for CSA  
Object deck for CSB  
OVERLAY ONE.  
Object deck for CSC  
OVERLAY TWO.  
Object deck for CSD  
Object deck for CSE  
OVERLAY TWO.  
Object deck for CSF  
OVERLAY ONE.  
Object deck for CSG  
/*
```

The EXEC statement illustrates that the OVLY parameter must be specified for every overlay program to be processed by the linkage editor.

Using INCLUDE Statements

The primary input data set for this example contains a series of control statements. The INCLUDE statements in the primary input data set direct the linkage editor to library members that contain the control sections of the program.

```
//LKED          EXEC  PGM=HEWL,PARM='OVLY'  
.  
.  
//MODLIB        DD    DSNAME=OBJLIB,DISP=(OLD,KEEP),...  
//SYSLIN        DD    *  
ENTRY BEGIN  
INCLUDE MODLIB(CSA,CSB)  
OVERLAY ONE  
INCLUDE MODLIB(CSC)  
OVERLAY TWO  
INCLUDE MODLIB(CSD,CSE)  
OVERLAY TWO  
INCLUDE MODLIB(CSF)  
OVERLAY ONE  
INCLUDE MODLIB(CSG)  
/*
```

This example differs from the previous one in that the control sections of the program are not part of the primary input data set, but are represented in the primary input by the INCLUDE statements. When an INCLUDE statement is processed, the appropriate control section is retrieved from the library and processed.

Using INSERT Statements

When INSERT statements are used, the INSERT and OVERLAY statements may either follow or precede all the input modules. However, the order of the control sections in a segment is not necessarily the same as the order of the INSERT statements for each segment. An example of each is given, as well as an example of repositioning automatically called control sections.

FOLLOWING ALL INPUT: The control statements can follow all the input modules, as shown in the following example:

```

//LKED      EXEC  PGM=HEWL,PARM='OVLY'
           .
           .
//SYSLIN    DD    DSNAME=OBJECT,DISP=(OLD,KEEP),...
//          DD    *
ENTRY BEGIN
INSERT CSA,CSB
OVERLAY ONE
INSERT CSC
OVERLAY TWO
INSERT CSD,CSE
OVERLAY TWO
INSERT CSF
OVERLAY ONE
INSERT CSG
/*

```

The primary input data set contains the object modules for the control sections, and the input stream is concatenated to it.

PRECEDING ALL INPUT: The control statements can also precede all input modules, as shown in the following example:

```

//LKED      EXEC  PGM=HEWL,PARM='OVLY'
//MODULES   DD    DSNAME=OBJSEQ,DISP=(OLD,KEEP),...
           .
           .
//SYSLIN    DD    *
ENTRY BEGIN
INSERT CSA,CSB
OVERLAY ONE
INSERT CSC
OVERLAY TWO
INSERT CSD,CSE
OVERLAY TWO
INSERT CSF
OVERLAY ONE
INSERT CSG
INCLUDE MODULES
/*

```

The primary input data set contains all the control statements for the overlay structure and an INCLUDE statement. The data set specified by the INCLUDE statement contains all the object modules for the structure, and is a sequential data set.

REPOSITIONING AUTOMATICALLY CALLED CONTROL SECTIONS: The INSERT statement can also be used to move automatically called control sections from the root segment to the desired segment. This is helpful when control sections from the automatic call library are used in only one segment. By moving such control sections, the root segment will contain only those control sections used by more than one segment.

When a program is written in a higher level language, special control sections are called from the automatic call library. Assume that the sample program is written in COBOL and that two control sections (ILBOVTR0 and ILBOSCH0) are called automatically from SYS1.COBLIB. Ordinarily, these control sections are placed in the root segment. However, INSERT statements are used in the following example to place these control sections in segments other than the root segment.

```
//LKED          EXEC  PGM=HEWL,PARM='OVLY'  
//MODLIB       DD    DSNAME=OBJLIB,DISP=(OLD,KEEP),...  
//SYSLIB       DD    DSNAME=SYS1.COBLIB,DISP=SHR  
.  
.  
//SYSLIN       DD    *  
  ENTRY BEGIN  
  INCLUDE MODLIB(CSA,CSB)  
  OVERLAY ONE  
  INCLUDE MODLIB(CSC)  
  OVERLAY TWO  
  INCLUDE MODLIB(CSD,CSE)  
  INSERT ILBOVTR0  
  OVERLAY TWO  
  INCLUDE MODLIB(CSF)  
  INSERT ILBOSCH0  
  OVERLAY ONE  
  INCLUDE MODLIB(CSG)  
/*
```

As a result, segments 3 and 4 will also contain ILBOVTR0 and ILBOSCH0, respectively.

This example also combines two of the ways of specifying the control sections for a segment.

SPECIAL OPTIONS

The linkage editor provides three special job step options (OVLY, LET, and XCAL) for the overlay programmer. These options are specified on the EXEC statement for the linkage editor job step. They must be respecified each time a load module in overlay structure is reprocessed by the linkage editor.

OVLY Option

The OVLY option must be specified for every overlay program. If the option is omitted, all the OVERLAY and INSERT statements are considered invalid, and the output module is not an overlay structure. If in addition, the LET option is not specified, the output module is marked not executable.

LET Option

With the LET option, the output module is marked executable even though certain error conditions were found during linkage editor processing. When LET is specified, any exclusive reference (valid or invalid) is accepted. At execution time, a valid exclusive reference is executed correctly; an invalid exclusive reference usually causes unpredictable results.

Also with the LET option, unresolved external references do not prevent the module from being marked executable. This could be helpful when part of a large program is ready for testing; the segments to be tested may contain references to segments not yet coded. If LET is specified, the program can be executed to test those parts that are finished (as long as the references to the absent segments are not executed). If the LET option is not specified, these unresolved references will cause the module to be marked not executable.

XCAL Option

With the XCAL option, a valid exclusive call is not considered an error, and the load module is marked executable. However, unless the LET option is specified, other errors could cause the module to be marked not executable. In this case, the XCAL option is not required.

AMODE and RMODE Options

If the OVLY option is specified, the AMODE and RMODE options are ignored and a diagnostic message is issued to that effect. Overlay programs are assigned a residence mode of 24 and an addressing mode of 24.

SPECIAL CONSIDERATIONS

This section discusses several special considerations that affect overlay programs. These considerations include the handling of common areas, special storage requirements, and overlay communication.

COMMON AREAS

When common areas (blank or named) are encountered in an overlay program, the common areas are collected as described previously (that is, the largest blank or identically named common area is used). The final location of the common area in the output module depends on whether INSERT statements were used to structure the program.

If INSERT statements are used to structure the overlay program, a named common area should either be part of the input stream in the segment to which it belongs, or should be placed there with an INSERT statement.

Because INSERT statements cannot be used for blank common areas, a blank common area should always be part of the input stream in the segment to which it belongs.

If INSERT statements are not used, and the control sections for each segment are placed or included between OVERLAY statements,

the linkage editor "promotes" the common area automatically. That is, the common area is placed in the common segment of the paths that contain references to it so that the common area is in storage when needed. The position of the promoted area in relation to other control sections within the common segment is unpredictable.

If a common area is encountered in a module from the automatic call library, automatic promotion places the common area in the root segment. In the case of a named common area, this may be overridden by use of the INSERT statement.

Assume that the sample program is written in FORTRAN and that common areas are present as shown in Figure 60. Further assume that the overlay program is structured with INCLUDE statements between the OVERLAY statements so that automatic promotion occurs.

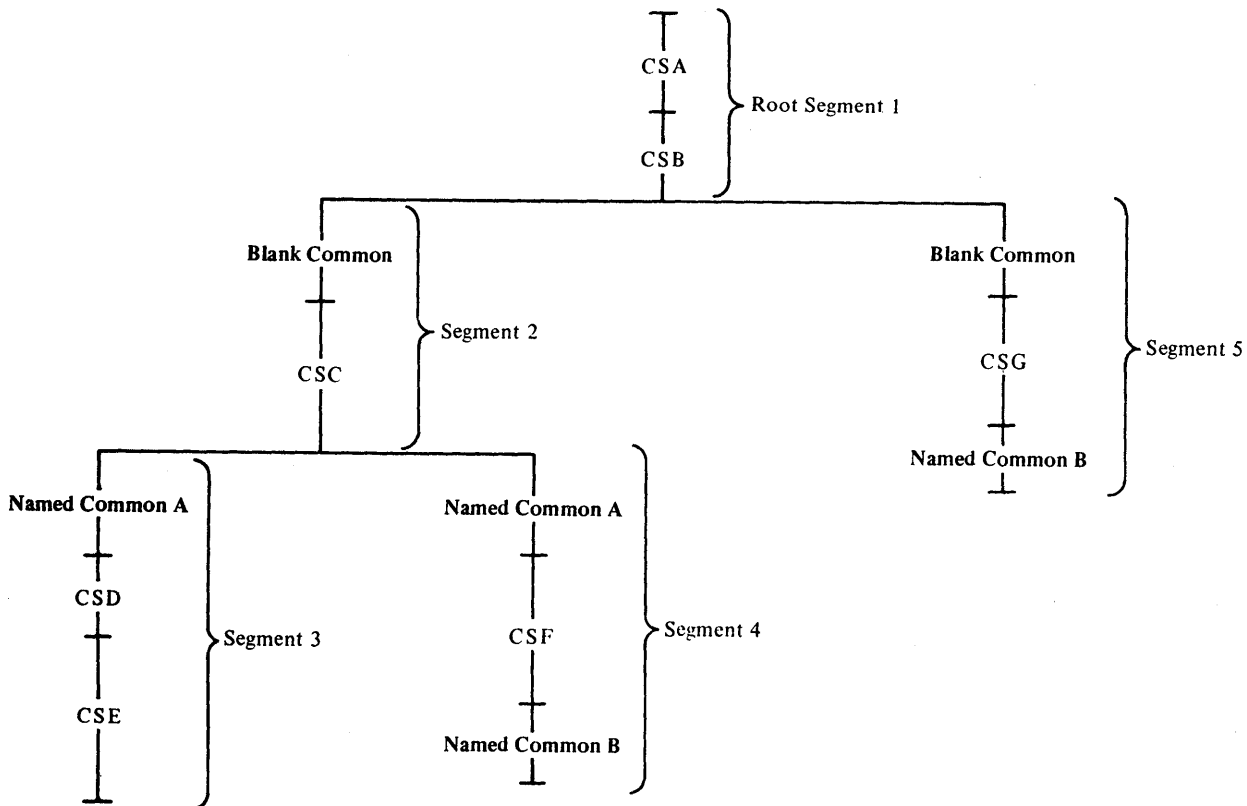


Figure 60. Common Areas before Processing

Segments 2 and 5 contain blank common areas, segments 3 and 4 contain named common area A, and segments 4 and 5 contain named common area B. During linkage editor processing, the blank common areas are collected and the largest area is promoted to the root segment (the first common segment in the two paths); the common areas named A are collected and the largest area is promoted to segment 2; the common areas named B are collected and promoted to the root segment. Figure 61 on page 159 shows the location of the common areas after processing by the linkage editor.

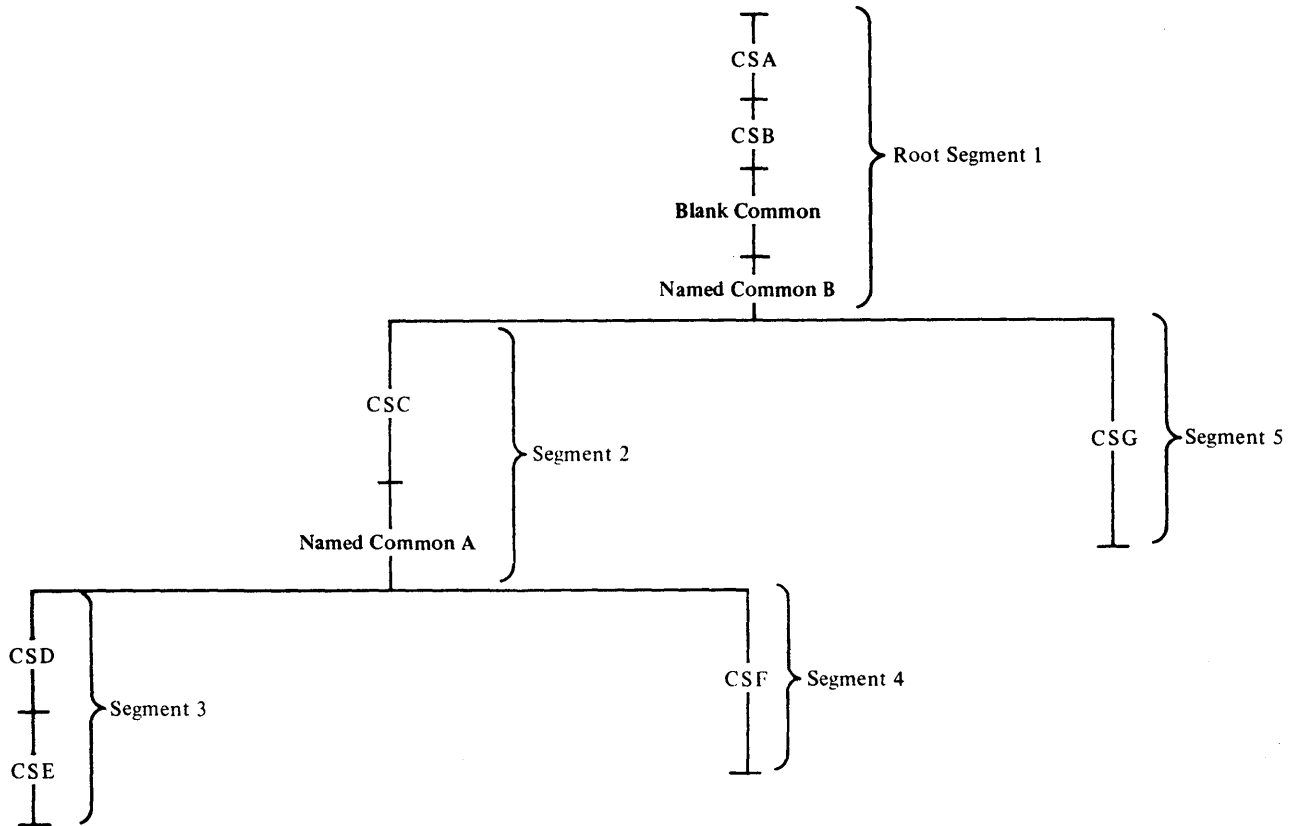


Figure 61. Common Areas after Processing

STORAGE REQUIREMENTS

The virtual storage requirements for an overlay program include the items placed in the module by the linkage editor and the overlay supervisor necessary for execution.

ITEMS IN THE LOAD MODULE: The items that the linkage editor places in an overlay load module are the segment table, entry tables, and other control information. Their size must be included in the minimum requirements for an overlay program, along with the storage required by the longest path and any control sections from the automatic call library.

Every overlay program has one segment table in the root segment. The storage requirements are:

$$\text{Length of SEGTAB} = (4n + 24) \text{ bytes}$$

where:

n = the number of segments in the program

Some segments will have an entry table. The requirements of the entry tables in the segments in the longest path must be added to the storage requirements for the program. The requirements for an entry table are:

Length of ENTAB = $12(x + 1)$ bytes

where:

x = the number of entries in the table

Finally, a NOTE list is required to execute an overlay program. The storage requirements are:

Length of NOTELST = $(4n + 8)$ bytes

where:

n = the number of segments in the program

OVERLAY SUPERVISOR: To the minimum requirements of the load module itself must be added the requirements of the overlay supervisor. This system routine is not placed in an overlay module, but, during execution of the module, the supervisor may be called to initiate an overlay. If called, the storage allocated for the program must also be large enough for the supervisor.

This asynchronous overlay supervisor module is furnished with the system. This asynchronous module also permits overlay through the SEGLD macro instruction (see "Overlay Communication"). The storage requirement for the overlay supervisor module is 180 bytes.

OVERLAY COMMUNICATION

Several ways of communicating between segments of an overlay program are discussed in this section. A higher level or Assembler language program may use a CALL statement or a CALL macro instruction, respectively, to cause control to be passed to a symbol defined in another segment. The CALL may cause the segment to be loaded if it is not already present in storage. An Assembler language program may also use three additional ways to communicate between segments:

- By a branch instruction, which causes a segment to be loaded and control to be passed to a symbol defined in that segment.
- By a segment load (SEGLD) macro instruction, which requests loading of a segment. Processing continues in the requesting segment while the requested segment is being loaded.
- By a segment load and wait (SEGWT) macro instruction, which requests loading of a segment. Processing continues in the requesting segment only after the requested segment is loaded.

Any of the four methods may be used to make inclusive references. Only the CALL and branch may be used to make exclusive references. Neither the SEGLD nor the SEGWT macro instruction should be used to make exclusive references; because both imply that processing is to continue in the requesting segment, an exclusive reference leads to erroneous results when the program is executed.

CALL Statement or CALL Macro Instruction

A CALL statement or a CALL macro instruction refers to an external name in the segment to which control is to be passed. The external name must be defined as an external reference in the requesting segment. In Assembler language, the name must be defined as a 4-byte V-type address constant; the high-order bit is reserved for use by the control program, and must not be altered during execution of the program.

When a CALL is used, the requested segment and any segments in its path are loaded if they are not part of the path already in virtual storage. After the segment is loaded, control is passed to the requested segment at the location specified by the external name.

A CALL between inclusive segments is always valid. A return can be made to the requesting segment by another source language statement, such as RETURN. A CALL between exclusive segments is valid if the conditions for a valid exclusive reference are met; a return from the requested segment can be made only by another exclusive reference, because the requesting segment has been overlaid.

Branch Instruction

Any of the branching conventions shown in Figure 62 can be used to request loading and branching to a segment. As a result, the requested segment and any segments in its path are loaded if they are not part of the path already in virtual storage. Control is then passed to the requested segment at the location specified by the address constant placed in general register 15.

Example	Name ¹	Operation	Operand ^{2,3}
1		L BALR	R15,=V(name) Rn,R15
2		L BALR	R15,ADCON Rn,R15
.		.	.
.	ADCON	DC	V(name)
3		L BAL	R15,=V(name) Rn,0(0,R15) ⁴
4		L BAL	R15,=V(name) Rn,0(R15) ⁵
5 ⁶		L BCR	R15,=V(name) 15,R15
6 ⁶		L BC	R15,=V(name) 15,0(0,R15) ⁴
7 ⁶		L BC	R15,=V(name) 15,0(R15) ⁵

Figure 62. Branch Sequences for Overlay Programs

Notes to Figure 62:

- ¹ When the name field is blank, specification of a name is optional.
- ² R15 must hold a 4-byte address constant that is the address of an entry name or a control section name in the requested segment. The address constant must be loaded into the standard entry point register, register 15.
- ³ Rn is any other register and is used to hold the return address. This register is usually register 14.
- ⁴ This may also be written so that the index register is loaded with the address constant; the other fields must be zero.
- ⁵ In this format, the base register must be loaded with the address constant; the displacement must be zero.
- ⁶ This example is an unconditional branch; other conditions are also allowed.

The address constant must be a 4-byte V-type address constant. The high-order bit is reserved for use by the control program, and must not be altered during execution of the program.

A branch between inclusive segments is always valid; a return may be made by means of the address stored in Rn. A branch between exclusive segments is valid if the conditions for a valid exclusive reference are met; a return can be made only by another exclusive reference.

Segment Load (SEGLD) Macro Instruction

The SEGLD macro instruction is used to provide overlap between segment loading and processing within the requesting segment. As a result of using any of the examples in Figure 63, the loading of the requested segment and any segments in its path is initiated when they are not part of the path already in virtual storage. Processing then resumes at the next sequential instruction in the requesting segment while the segment or segments are being loaded. Control may be passed to the requested segment with either a CALL or a branch, as shown in Examples 1 and 2, respectively. A SEGWT instruction can be used to ensure that the data in the control section specified by the external name is in virtual storage before processing resumes, as shown in Example 3.

Example	Name ¹	Operation	Operand ^{2,3}
1		SEGLD CALL	external name external name
2		SEGLD branch	external name external name
3		SEGLD SEGWT L	external name external name Rn,=V(name)

Figure 63. Use of the SEGLD Macro Instruction

Notes to Figure 63:

- ¹ When the name field is blank, specification of a name is optional.
- ² External name is an entry name or a control section name in the requested segment.
- ³ Rn is any other register and is used to hold the return address. This register is usually register 14.

The external name specified in the SEGLD macro instruction is defined with a 4-byte V-type address constant. The high-order bit is reserved for use by the control program and must not be altered during execution of the program.

Segment Wait (SEGWT) Macro Instruction

The SEGWT macro instruction is used to stop processing in the requesting segment until the requested segment is in virtual storage.

As a result of using any of the examples in Figure 64, no further processing takes place until the requested segment and all segments in its path are loaded when not already in virtual storage. Processing resumes at the next sequential instruction in the requesting segment after the requested segment has been loaded.

Example	Name ¹	Operation	Operand ^{2,3}
1		SEGLD	external name
		SEGWT L	external name Rn,ADCON
	ADCON	branch DC	V(name)
2		SEGWT L	external name Rn,=V(name)

Figure 64. Use of the SEGWT Macro Instruction

Notes to Figure 64:

- ¹ When the name field is blank, specification of a name is optional.
- ² External name is an entry name or a control section name in the requested statement.
- ³ Rn is any other register and is used to hold the return address. This register is usually register 14.

If the SEGWT and SEGLD macro instructions are used together, overlap occurs between processing and segment loading; use of the SEGWT macro instruction serves as a check to see that the necessary information is in storage when it is finally needed (see Example 1 in Figure 64). In Example 2 in Figure 64, no overlap is provided; the SEGWT macro instruction initiates loading, and processing is stopped in the requesting segment until the requested segment is in virtual storage.

The external name specified in the SEGWT macro instruction must be defined with a 4-byte V-type address constant. The high-order bit is reserved for use by the control program, and must not be altered during execution of the program.

If the contents of a virtual storage location in the requested segment are to be processed, the entry name of the location must be referred to by an A-type address constant.

PART II. LOADER

CHAPTER 9. OVERVIEW AND USES OF THE LOADER

The loader is a processing program that combines basic editing and loading functions of the linkage editor and program fetch into one job step. Therefore, the load function is equivalent to the link-edit-go function. The loader can be used for compile-load and load jobs.

The loader will load object modules produced by a language processor and load modules produced by the linkage editor into virtual storage for execution. Optionally, it will search a call library (SYSLIB) or a resident link pack area, or both, to resolve external references. The loader does not produce load modules for program libraries.

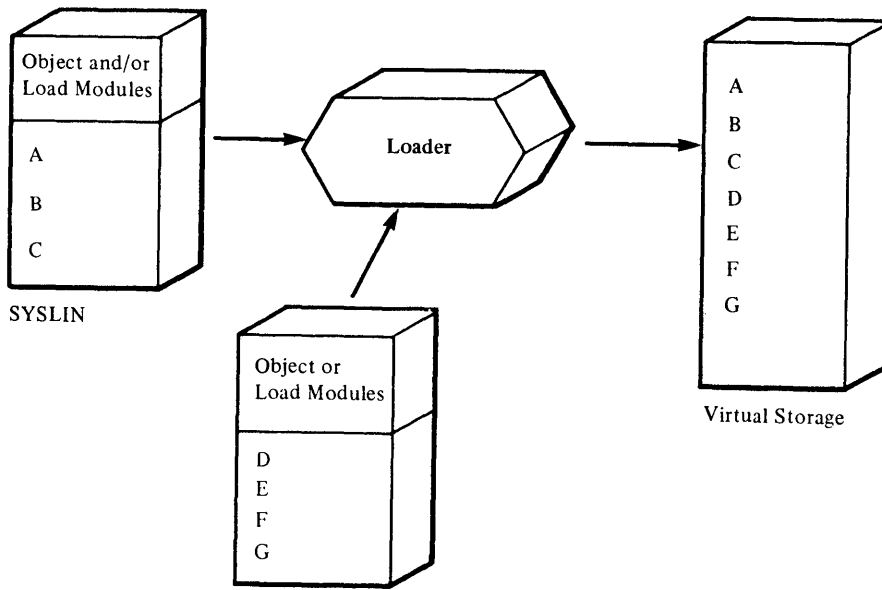
The functional characteristics, compatibility and restrictions, performance considerations, and storage considerations of the loader are described in the following sections.

FUNCTIONAL CHARACTERISTICS

The loader is reenterable and, therefore, can reside in the resident link pack area.

The loader combines the following basic functions of the linkage editor and program fetch:

1. Resolution of external references between program modules.
2. Optional inclusion of modules from a call library (SYSLIB) or from a link pack area, or from both (Figure 65 on page 167 and Figure 66 on page 168). (Inclusion of modules from a call library or the link pack area is performed, if requested, when external references remain unresolved after processing the primary input to the loader. If both are requested, the link pack area is searched first.)
3. Automatic deletion of duplicate copies of program modules (Figure 67 on page 168). (The first copy is loaded and all following requests use that copy.)
4. Relocation of all address constants so that control may be passed directly to the assigned entry point in virtual storage.



SYSLIB – called automatically when references were unresolved at the end of input from SYSLIN.

Figure 65. Loader Processing—SYSLIB Resolution

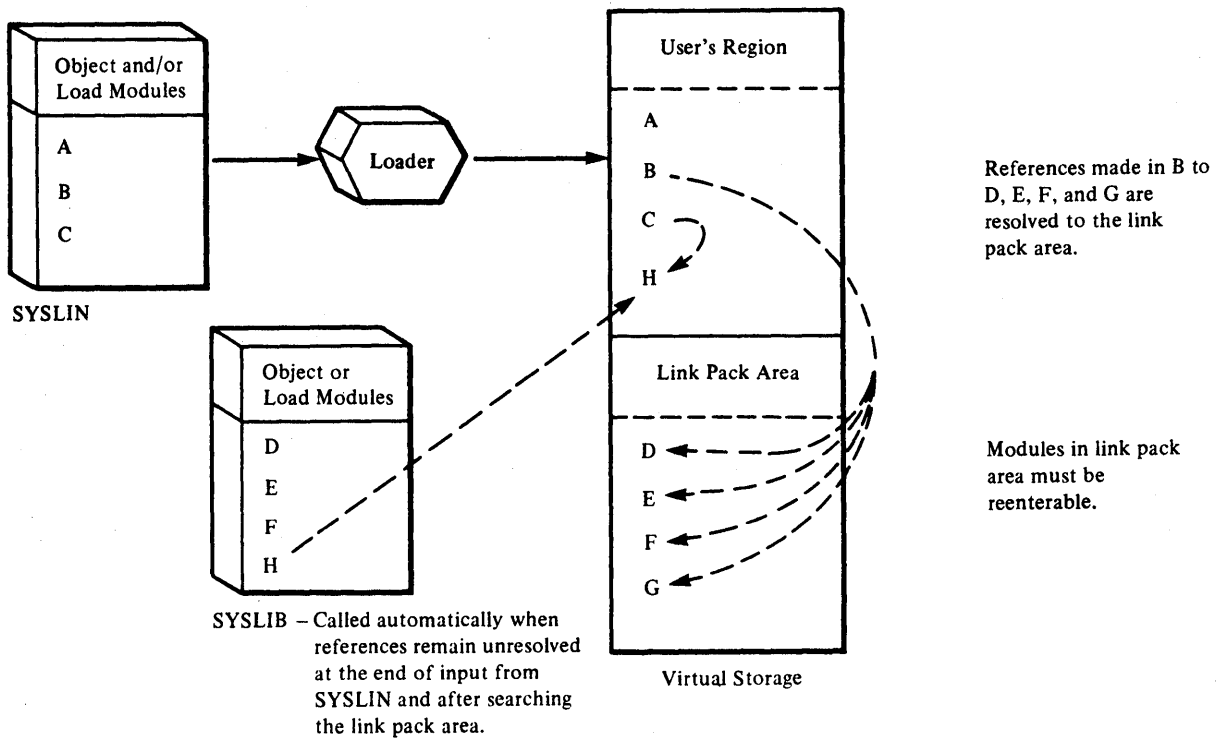


Figure 66. Loader Processing—Link Pack Area and SYSLIB Resolution




Figure 67. Loader Processing—Automatic Editing

COMPATIBILITY AND RESTRICTIONS

The loader accepts the same basic input as the linkage editor:

1. All object modules that can be processed by the linkage editor can be input to the loader.
2. All load modules produced by the linkage editor can be input to the loader (except load modules edited with the NE option).

The loader supports the following linkage editor options: MAP, LET, NCAL, SIZE, and TERM. All other linkage editor options and attributes are not supported, but, if used, they will not be considered as errors. A message will be listed on SYSLOUT indicating that they are not supported. The supported options are specified in the PARM field of the EXEC statement, or with the LINK, ATTACH, LOAD, or XCTL macro instruction. In addition to the supported linkage editor options, the loader provides several other options. All loader options are described under "EXEC Statement" on page 170.

The loader does not process linkage editor control statements (for example, INCLUDE, NAME, and OVERLAY). If they are used, they will not be treated as errors, and a message will be listed on SYSLOUT indicating that the control statements are not supported.

The loader and the linkage editor are bound by the same input conventions. (These conventions are discussed in Part 1 of this publication.) In addition, the loader can accept load modules in the SYSLIN data set and object modules from a data area in virtual storage.

The loader does not use auxiliary storage space for work areas; that is, there is no loader function corresponding to the linkage editor's creation of intermediate work data sets or output load modules.

Time Sharing Option (TSO)

When the loader is used under TSO, it is invoked by the loader prompter, a program that acts as an interface between the user and the operating system and the loader. Under TSO, execution of the loader and definition of the data sets used by the loader are described to the system through use of the LOADGO command that causes the prompter to be executed. Operands of the LOADGO command can also be used to specify the loader options a job requires.

Complete procedures for using the LOADGO command to load and execute an object module are given in TSO Command Language Reference.

Processing Object Modules in Virtual Storage

The loader can act as an interface with a compiler that has the ability to construct a data area of one or more object modules in virtual storage as an alternative to a data set on a secondary storage volume (such as a tape or disk). The compiler passes the loader a description of the internal data area, which the loader then processes as primary input. This internal data area replaces external SYSLIN data set input to the loader.

Instead of placing text records for the object module in the internal data area, the compiler can pass pointers to preloaded text. The loader can then perform its relocation and linkage functions on the preloaded text itself; text is not moved during processing.

CHAPTER 10. PREPARING INPUT FOR THE LOADER

This chapter discusses how to prepare an input deck for the loader and how to invoke the loader; it also describes the output from the loader.

INPUT FOR THE LOADER

The input deck for the loader must contain job control language statements for the loader and, optionally, for the loaded program (Figure 68).

```
//name      JOB      parameters      (optional)
//name      EXEC     PGM=LOADER,
              PARM=(parameters)
//SYSLIN    DD       parameters
//SYSLIB    DD       parameters      (optional)
//SYSLOUT   DD       parameters      (optional)
//SYSTEM    DD       parameters      (optional)

//          (optional DD statements and data
//          required for loaded program)
```

Figure 68. Input Deck for the Loader—Basic Format

Only the EXEC statement and the SYSLIN DD statement are required for a loader step. The JOB statement is required if the loader is the first step in the job.

EXEC STATEMENT

The EXEC statement is used to call the loader and to specify options for the loader and the loaded program. The loader is called by specifying PGM=HEWLDRGO or PGM=LOADER (see "Chapter 11. Invoking The Loader" on page 180).

PARM Field Format

Loader and loaded program options are specified in the PARM field of the EXEC statement. The PARM field must have the following format:

```
,PARM='(loaderoption(...))(/programoption(...))'
```

Note that the loaded program option(s), if any, must be separated from the loader option(s) by a slash (/). If there are no loader options, the program option must begin with a slash. The entire PARM field may be omitted if there are no loader or loaded program options.

Parameters must be enclosed in single quotation marks when special characters (/ and =) are used.

LOADER OPTIONS

The loader options are outlined below. Fields that must be supplied by the user are underlined; default options are underlined in bold-faced type.

Parameter	Options
PARM=	<u>CALL</u> <u>NOCALL</u> EP= <u>name</u> LET <u>NOLET</u> MAP <u>NOMAP</u> NAME= <u>name</u> <u>NAME=**GO</u> <u>PRINT</u> <u>NOPRINT</u> <u>RES</u> <u>NORES</u> SIZE= <u>size</u> <u>SIZE=300K</u> TERM <u>NOTERM</u>

CALL|NOCALL: Automatically Searching SYSLIB

Explanation: CALL|NOCALL are options specifying whether an automatic search of the SYSLIB data set is made when the loader is invoked.

CALL

Indicates that the SYSLIB data set will automatically be searched when the loader is invoked.

NOCALL

Indicates that no automatic search of the SYSLIB data set is to be made.

Abbreviations:

NOCALL|NCAL

Default: The default is CALL.

Restrictions: If the SYSLIB DD statement is not included in the input deck, the CALL|NOCALL option is ignored.

If the NOCALL option is specified, the NORES option is automatically set.

EP=name: Specifying the Program Entry Point

Explanation: EP=name is a loader option that specifies the external name to be assigned as the entry point of the loaded program.

Abbreviations: None.

Default: None.

Restrictions: EP=name must be specified if the entry point of the loaded program is in an input load module.

For FORTRAN, ALGOL, and PL/I, these entry points must be named MAIN, IHIFSAIN, and IHENTRY, respectively, unless changed by compiler options.

LET|NOLET: Executing with Severity 2 Errors

Explanation: LET|NOLET are loader options specifying whether the loader should try to execute the object program if a severity 2 error condition is found. A severity 2 error condition is one that could make execution of the loaded program impossible.

LET

Indicates that the loader will attempt to execute the program even if a severity 2 error is found.

NOLET

Indicates that the loader will not attempt to execute the program if a severity 2 error is found.

Abbreviations: None.

Default: The default is NOLET.

MAP|NOMAP: Printing a Program Map

Explanation: MAP|NOMAP are loader options specifying whether to produce a map of the loaded program that lists external names and their absolute storage addresses on the SYSLOUT data set. The module map is described in "Chapter 12. Interpreting Loader Output" on page 185.

MAP

Indicates that a program map will be printed.

NOMAP

Indicates that a program map will not be printed.

Abbreviations: None.

Default: The default is NOMAP.

Restrictions: If the SYSLOUT DD statement is not used in the input deck, the MAP|NOMAP option is ignored.

NAME=name: Identifying the Loaded Program

Explanation: NAME=name is a loader option specifying the name to be used to identify the loaded program to the system.

Abbreviations: None.

Default: If this option is not used, the loaded program will be named *XGO.

PRINT|NOPRINT: Printing Messages on SYSLOUT

Explanation: PRINT|NOPRINT are loader options specifying that informational and diagnostic messages are to be produced on the SYSLOUT data set.

PRINT

Indicates that messages are printed in SYSLOUT.

NOPRINT

Indicates that no messages are printed in SYSLOUT.

Abbreviations: None.

Default: The default is PRINT.

Restrictions: If NOPRINT is specified, the SYSLOUT data set is not opened.

RES|NORES: Automatically Searching the Link Pack Area Queue

Explanation: RES|NORES are loader options specifying whether an automatic search of the link pack area queue is to be made after processing the primary input (SYSLIN) and before searching the SYSLIB data set.

RES

Indicates that an automatic search of the link pack area queue is to be made.

NORES

Indicates that no automatic search of the link pack area queue is to be made.

Abbreviations: None.

Default: The default is RES.

Restrictions: When the RES option is specified, the CALL option is also automatically set.

SIZE=size: Specifying Virtual Storage

Explanation: SIZE=size is a loader option specifying the amount of dynamic virtual storage, in bytes, that can be used by the loader. See "Appendix E. Loader Return Codes" on page 189 for more information on size.

Abbreviations: None.

Default: If this option is not specified, the size defaults to 300K bytes.

TERM|NOTERM: Sending Messages to SYSTEM

Explanation: TERM|NOTERM are loader options specifying whether to send numbered diagnostic messages to the SYSTEM data set. Although TERM|NOTERM is intended to be used when operating under the Time Sharing Option (TSO), the SYSTEM data set can be used to replace or supplement the SYSLOUT data set at any time.

TERM

Indicates that numbered diagnostic messages are sent to the SYSTEM data set.

NOTERM

Indicates that no messages are to be sent to SYSTEM.

Abbreviations: None.

Default: The default is NOTERM.

Restrictions: If the SYSTEM DD statement is not included in the input deck, the TERM option is ignored.

EXEC STATEMENT EXAMPLE

The following are examples of the EXEC statement. In these examples, X and Y are parameters required by the loaded program.

```

//LOAD      EXEC    PGM=LOADER
//LOAD      EXEC    PGM=HEWLDRGO,
//          PARM='MAP,EP=FIRST/X,Y'
//LOAD      EXEC    PGM=LOADER,PARM='/X,Y'
//LOAD      EXEC    PGM=LOADER,PARM=NOPRINT
//LOAD      EXEC    PGM=LOADER,PARM=(MAP,LET)
//LOAD      EXEC    PGM=LOADER,
//          PARM='NAME=NEWPROG,TERM,NOPRINT'
//

```

For further details in coding the EXEC statement, refer to the publication JCL.

DD STATEMENTS

The loader uses four DD statements, named SYSLIN, SYSLIB, SYSLOUT, and SYSTEMR. The SYSLIN DD statement must be used in every loader job. The others are optional.

The following considerations apply to the DCB parameter of SYSLIN, SYSLIB, and SYSLOUT.

- For better performance, BLKSIZE and BUFNO can be specified.
- If BUFNO is omitted, BUFNO=2 is assumed.
- Any value given to BUFNO is assumed for NCP (number of channel programs).
- If RECFM=U is specified, BUFNO=2 is assumed, and BLKSIZE and LRECL are ignored.
- RECFM=V is not accepted.
- RECFM=FBSA is always assumed for SYSLOUT.
- If RECFM is omitted, RECFM=F is assumed for SYSLIN and SYSLIB.
- If BLKSIZE is omitted, the value given to LRECL is assumed.
- LRECL=121 is assumed for SYSLOUT unless the loader is operating under the Time Sharing Option (TSO), when LRECL=81 is assumed.
- If LRECL is omitted, LRECL=80 is assumed for SYSLIN and SYSLIB.
- If OPTCD=C is used to specify chained scheduling, and if the necessary data management routines are not resident, an additional 2K bytes (2048 bytes) of virtual storage is needed in the user's region.

Note: The SYSTEMR data set will always consist of unblocked 81-character records with BUFNO=2 and RECFM=FSA. Because these values are fixed, the DCB parameter need not be used. ihl.RECFM (record format)

In addition to the DD statements used by the loader, any DD statements and data required by the loaded program must be included in the input deck.

SYSLIN DD Statement

The SYSLIN DD statement defines the input data for the loader. This input can be either object modules produced by a language translator, or load modules produced by the linkage editor, or both. The data sets defined by the SYSLIN DD statement can be either sequential data sets or members of a partitioned data set, or both. The DSNAMES parameter for a partitioned data set must indicate the member name, that is,

```
DSNAME=dsname(membername).
```

Concatenation can be used to include more than one module in SYSLIN.

The following are examples of the SYSLIN DD statement. The first example defines a member of a previously cataloged partitioned data set:

```
//SYSLIN      DD      DSNAMES=OUTPUT.FORT(MOD12),  
//                               DISP=OLD,DCB=BLKSIZE=3200
```

The second example defines a sequential data set on magnetic tape:

```
//SYSLIN      DD      DSNAMES=PROG15,UNIT=3400-6,DISP=(OLD,  
//                               KEEP),VOLUME=(PRIVATE,RETAIN,  
//                               SER=MCS167)
```

The third example defines a data set that was the output of a previous step in the same job:

```
//SYSLIN      DD      DSNAMES=*.COBOL.SYSLIN,DISP=(OLD,  
//                               DELETE)
```

The fourth example shows the concatenation of three data sets. The first two data sets are members of different partitioned data sets; the first is an object module, and the second is a load module. The third data set is in the input stream following a SYSLIN DD statement (see "Loaded Program Data" on page 177).

```
//SYSLIN      DD      DSNAMES=PGMLIB.SET1(RFS1),DISP=OLD,  
//                               DCB=(BLKSIZE=3200,RECFM=FB)  
//                               DD      DSNAMES=PGMLIB.SET2(ABC5),DISP=OLD,  
//                               DCB=RECFM=U  
//                               DD      DDNAME=SYSIN
```

SYSLIB DD Statement

The SYSLIB data set contains IBM-supplied or user-written library routines to be included in the loaded program. The data set is searched when unresolved references remain after processing SYSLIN and optionally searching the link pack area.

The SYSLIB data set is used to resolve an external reference when the following conditions exist: the external reference must be (1) a member name or an alias of a module in the data set, and (2) defined as an external name in the external symbol dictionary of the module with that name. If the unresolved external reference is a member name or an alias in the library, but is not an external name in that member, the member is processed but the external reference remains unresolved unless subsequently defined.

The data set defined by the SYSLIB DD statement must be a partitioned data set that contains either object modules or load modules, but not both. Concatenation may be used to include more partitioned data sets in SYSLIB. All concatenated data sets must contain the same type of modules (object or load).

The following are examples of the SYSLIB DD statement. The first example defines a cataloged partitioned data set that can be shared by other steps:

```
//SYSLIB DD DSNAME=SYS1.ALGLIB,DISP=SHR
```

The second example shows the concatenation of two data sets:

```
//SYSLIB DD DSNAME=SYS1.PL1LIB,DISP=SHR
// DD DSNAME=LIBMOD.MATH,DISP=OLD
```

SYSLOUT DD Statement

The SYSLOUT DD statement is used for error and warning messages and for an optional map of external references (see "Chapter 12. Interpreting Loader Output" on page 185). The data set defined by this DD statement must be a sequential data set. The DCB parameter can be used to specify the blocking factor (BLKSIZE) of this data set. For better performance, the number of buffers (BUFNO) to be allocated to SYSLOUT can also be specified.

The following are examples of the SYSLOUT DD statement. The first example specifies the system output unit:

```
//SYSLOUT DD SYSOUT=A
```

The second example defines a sequential data set on a 3800 printer:

```
//SYSLOUT DD UNIT=3800,DCB=(BLKSIZE=121,
// BUFNO=4)
```

SYSTEM DD Statement

The SYSTEM DD statement defines a data set that is used for numbered diagnostic messages only. When the loader is being used under the Time Sharing Option (TSO) of the operating system, the SYSTEM DD statement defines the terminal output data set. However, SYSTEM can also be used at any time to replace or supplement the SYSLOUT data set. Because the SYSTEM data set is not opened unless the loader must issue a diagnostic message, using SYSTEM instead of SYSLOUT can reduce loader processing time.

When the SYSTEM data set replaces the SYSLOUT data set, the numbered messages in the SYSTEM data set are the only diagnostic output; when SYSTEM supplements the SYSLOUT data set, the numbered messages appear in both data sets, and optional diagnostic and informational output, such as a list of options or a module map, can be obtained on SYSLOUT.

The DCB parameters for SYSTEM are fixed and need not be specified. The SYSTEM data set always consists of unblocked 81-character records with BUFNO=2 and RECFM=FSA.

The following example shows the SYSTEM DD statement when used to specify the system output unit:

```
//SYSTEM DD SYSOUT=A
```

LOADED PROGRAM DATA

Loaded program data and loader data can both be specified in the input reader. Loaded program data can be defined by a DD statement following the loader data.

Figure 69 shows the loading of a previously compiled FORTRAN problem program. The program to be loaded (loader data) follows the SYSLIN DD statement. The loaded program data follows the FT05F001 DD statement.

```
//LOAD      JOB      MSGLEVEL=1
//LDR       EXEC     PGM=LOADER, PARM=MAP
//SYSLIB    DD       DSNAME=SYS1.FORTLIB, DISP=SHR
//SYSLOUT   DD       SYSOUT=A
//FT06F001  DD       SYSOUT=A
//SYSLIN    DD       *

                (Loader data)

/*
//FT05F001  DD       *

                (Loaded program data)

/*
```

Figure 69. Loader and Loaded Program Data Input Stream

SAMPLE INPUT FOR THE LOADER

Figure 70 shows an input deck for a load job. A previously assembled program, MASTER, is to be loaded. The SYSLOUT, SYSLIB, and SYSTEM DD statements are not used.

```
//LOAD      JOB      MSGLEVEL=1
//         EXEC     PGM=LOADER
//SYSLIN    DD       DSNAME=MASTER, DISP=OLD

                (DD statements and data required for execution of MASTER)

/*
```

Figure 70. Input Deck for a Load Job

Figure 71 shows an input deck for a compile-load job. The OS/VS COBOL (IKFCBL00) compiler is used for the compile step. The loaded program requires the SYSOUT, TAXRATE, and SYSIN DD statements.

```
//JOB          JOB      22,MCS,MSGLEVEL=1
//COBOL        EXEC     PGM=IKFCBL00,PARM=DMAP,REGION=256K,RD=R
//SYSPRINT     DD       SYSOUT=A
//SYSPUNCH     DD       SYSOUT=B
//SYSUT1       DD       UNIT=SYSDA,SPACE=(TRK,(100,10))
//SYSUT2       DD       UNIT=SYSDA,SPACE=(TRK,(100,10))
//SYSUT3       DD       UNIT=SYSDA,SPACE=(TRK,(100,10))
//SYSUT4       DD       UNIT=SYSDA,SPACE=(TRK,(100,10))
//SYSLIN       DD       DSNAME=##LOADSET,DISP=(MOD,PASS),
//              DD       UNIT=SYSSQ,SPACE=(TRK,(30,10))
//SYSIN        DD       *
```

(source program)

```
//LOAD         EXEC     PGM=LOADER,PARM='MAP,LET',COND=(5,LT,
//              DD       COBOL)
//SYSLIN       DD       DSNAME=*.COBOL.SYSLIN,DISP=(OLD,
//              DD       DELETE)
//SYSLOUT      DD       SYSOUT=A
//SYSLIB       DD       DSNAME=SYS1.COBLIB,DISP=SHR
//SYSOUT       DD       SYSOUT=A
//TAXRATE      DD       DSNAME=TAXRATE,DISP=OLD
//SYSIN        DD       *
```

(Data for Loaded Program)

/*

Figure 71. Input Deck for a Compile-Load Job

Figure 72 on page 179 shows the compilation and loading of three modules. In the first three steps, the OS/VS FORTRAN (FORTVS) compiler is used to compile a main program, MAIN, and two subprograms, SUB1 and SUB2. Each of the object modules is placed in a sequential data set by the compiler and passed to the loader job step. In addition to the FORTRAN library, a private library, MYLIB, is used to resolve external references. In the loader job step, MYLIB is concatenated with the SYSLIB DD statement. SUB1 and SUB2 are included in the program to be loaded by concatenating them with the SYSLIN DD statement. The SYSTEM statement is used to define the diagnostic output data set. The loaded program requires the FT01F001 and FT10F001 DD statements for execution, and it does not require data in the input stream.

```

//JOBX      JOB
//STEP1     EXEC PGM=FORTVS, PARM='NAME=MAIN, LOAD'
           .
           .
//SYSLIN    DD  DSNAME=&&GOFIL, DISP=(,PASS), UNIT=SYSSQ
//SYSIN     DD  *

      (Source module for MAIN)

/*
//STEP2     EXEC PGM=FORTVS, PARM='NAME=SUB1, LOAD'
           .
           .
//SYSLIN    DD  DSNAME=&&SUBPROG1, DISP=(,PASS), UNIT=SYSSQ
//SYSIN     DD  *

      (Source module for SUB1)

/*
//STEP3     EXEC PGM=FORTVS, PARM='NAME=SUB2, LOAD'
           .
           .
//SYSLIN    DD  DSNAME=&&SUBPROG2, DISP=(,PASS), UNIT=SYSSQ
//SYSIN     DD  *

      (Source module for SUB2)

/*
//STEP4     EXEC PGM=LOADER
//SYSTEM    DD  SYSOUT=A
//SYSLIB    DD  DSNAME=SYS1.FORTLIB, DISP=OLD
//          DD  DSNAME=MYLIB, DISP=OLD
//SYSLIN    DD  DSNAME=*.STEP1.SYSLIN, DISP=OLD
//          DD  DSNAME=*.STEP2.SYSLIN, DISP=OLD
//          DD  DSNAME=*.STEP3.SYSLIN, DISP=OLD
//FT01F001  DD  DSNAME=PARAMS, DISP=OLD
//FT10F001  DD  SYSOUT=A
/*

```

Figure 72. Input Deck for Compilation and Loading of the Three Modules

CHAPTER 11. INVOKING THE LOADER

The loader can be referred to by either its program name, HEWLDGR0, or its alias, LOADER. The loader can be invoked through the EXEC statement, as described in "Input for the Loader" on page 170, or through one of the following macro instructions.

[symbol]	LINK	EP= <u>loadername</u> , PARAM=(<u>optionlist</u> [, <u>ddname list</u>]), VL=1
[symbol]	ATTACH	EP= <u>loadername</u> , PARAM=(<u>optionlist</u> [, <u>ddname list</u>]), VL=1
[symbol]	LOAD	EP= <u>loadername</u>
[symbol]	XCTL	EP= <u>loadername</u>

EP=loadername

specifies the symbolic name of the loader. The entry point at which execution is to begin is determined by the control program from the library directory entry.

PARAM=(optionlist[,ddname list])

specifies, as a sublist, address parameters to be passed to the loader. The first fullword in the address parameter list contains the address of the option list for the loader and/or loaded program. The second fullword contains the address of the ddname list. If standard ddnames are to be used, ddname list may be omitted.

optionlist

specifies the address of a variable-length list containing the loader and loaded program options. This address must be written even though no real list is provided; for example, when the optionlist points to a halfword of zero.

The option list must begin on a halfword boundary. The two high-order bytes contain a count of the number of bytes in the remainder of the list. If no options are specified, the count must be zero.

The option list is free form, with the loader and loaded program options separated by a slash (/), and with each option separated by a comma. No blanks or zeros should appear in the list.

ddname list

specifies the address of a variable-length list containing alternative ddnames for the data sets used during loader processing. If the standard ddnames are used, this operand may be omitted.

The format of the ddname list is identical to the format of the ddname list for invoking the linkage editor; the 8-byte entries in the list are as follows:

Entry	Alternate Name For:
1	SYSLIN
2	not applicable
3	not applicable
4	SYSLIB
5	not applicable
6	SYSLOUT
7-11	not applicable
12	SYSTEM

VL=1

specifies that the sign bit is to be set to 1 in the last fullword of the address parameter list.

Figure 73 shows an Assembler language program that uses the LINK macro instruction to refer to the loader.

```

                SAVE      (14,12)                Initialize save
                .                    registers and point
                .                    to new save area
                LA        13,SAVEAREA
                .
                LINK      EP=LOADER,PARAM=(PARM),VL=1
                .
                .
                L         13,4(13)
                RETURN    (14,12),T
                .
                .
                DS        OH
PARM            DC        AL2(LENGTH)           Length of options
OPTIONS        DC        C'NOPRINT,CALL/X,Y,Z' loader and loaded
LENGTH        EQU      *-OPTIONS             program options
SAVEAREA      DS        18F                   Save area
                .
                .
                END

```

Figure 73. Using the LINK Macro Instruction to Refer to the Loader

If desired, the loader may be used to process a program but not execute it. To invoke just the portion of the loader that processes input data, specify either the name HEWLOAD or the name HEWLOADR with a LOAD and CALL macro instruction.

HEWLOAD loads and identifies the program. HEWLOAD returns the address of an 8-character name in register 1. This name can be used with an ATTACH, LINK, LOAD, or XCTL macro instruction to invoke the loaded program. A user program that is going to attach a loaded program should avoid specifying SZERO=NO in its ATTACH macro. If SZERO=NO must be specified, the user program should issue a LOAD for the loaded program before performing the ATTACH and a DELETE for the loaded program after the ATTACH.

HEWLOADR loads the program but will not identify it. HEWLOADR returns the entry point of the loaded program in register 0. Register 1 points to two fullwords: the first points to the beginning of storage occupied by the loaded program; the second contains the size of the loaded program. This location and size can then be used in a FREEMAIN macro instruction to free the storage occupied by the loaded program when it is no longer needed.

Figure 74 on page 183 shows an Assembler language program that uses the LOAD and CALL macro instructions to refer to HEWLOADR. Figure 75 on page 184 shows an Assembler language program that uses the LOAD and CALL macro instructions to refer to HEWLOAD.

```

SAVE      (14,12),T          Initialize save registers
        .                    and point to new save area
        .
        .
ST        13,SAVEAREA+4
LA        13,SAVEAREA
        .
        .
LOAD      EP=HEWLOADR        Load the loader
LR        15,0              Get its entry point address
CALL      (15),(PARM1),VL=1  Invoke the loader
        .
        .
LR        7,15              Save return code
LR        5,0              Save entry to loaded program
LR        6,1              Save pointer to list containing
*                                     Start address and length
DELETE    EP=HEWLOADR        Delete loader
CH        7,-H'4'          Verify successful loading
BH        FREE             Negative branch
LR        15,5            Loading successful—get entry
*                                     point address for CALL
CALL      (15),(PARM2),VL=1  Invoke program
        .
        .
FREE      L      0,4(6)      Get length into register 0
        L      1,0(6)      Get start address
FREEMAIN  R,LV=(0),A=(1)    Delete loaded program
        .
        .
L        13,4(13)
RETURN   (14,12),T
DS       OH
PARM1    DC      AL2(LENGTH1)  Length of loader options
OPTIONS1 DC      C'NOPRINT,CALL' Loader options
LENGTH1  EQU     *-OPTIONS1
        DS       OH
PARM2    DC      AL2(LENGTH2)  Length of loaded program options
OPTIONS2 DC      C'X,Y,Z'     Loaded program options
LENGTH2  EQU     *-OPTIONS2
SAVEAREA DS      18F         Save area
        .
        .
END

```

Figure 74. Using the LOAD and CALL Macro Instructions to Refer to HEWLOADR (Loading Without Identification)

	SAVE	(14,12),T	Initialize save registers and point to new save area
	.		
	.		
	ST	13,SAVEAREA+4	
	LA	13,SAVEAREA	
	.		
	.		
	LOAD	EP=HEWLOAD	Load the loader
	LR	15,0	Get its entry point address
	CALL	(15),(PARM1),VL=1	Invoke the loader
	LR	7,15	Save the return code
	MVC	PGMNAM(8),0(1)	Save program name
	DELETE	EP=HEWLOAD	Delete the loader
	CH	7,=H'4'	Verify successful loading
	BH	ERROR	Negative branch
	LINK	EPLOC=PGMNAM,PARM=(PARM2),VL=1	
	.		
	.		Loading successful, invoke program
ERROR	L	13,4(13)	
	RETURN	(14,12),T	
	DS	0H	
PARM1	DC	AL2(LENGTH1)	Length of loader options
OPTIONS1	DC	C'MAP'	Loader options
LENGTH1	EQU	*-OPTIONS1	
	DS	0H	
PARM2	DC	AL2(LENGTH2)	Length of loaded program options
OPTIONS2	DC	C'X,Y,Z'	Loaded program options
LENGTH2	EQU	*-OPTIONS2	
SAVEAREA	DS	18F	Save area
PGMNAM	DS	2F	Program name
	.		
	.		
	END		

Figure 75. Using the LOAD and CALL Macro Instructions to Refer to HEWLOAD (Loading With Identification)

For further information on the use of these macro instructions, see Supervisor Services and Macro Instructions.

CHAPTER 12. INTERPRETING LOADER OUTPUT

Loader output consists of a collection of diagnostics and error messages, and of an optional storage map of the loaded program. This output is produced in the data set defined by the SYSLOUT DD and SYSTEM DD statements. If these are omitted, no loader output is produced.

SYSLOUT output includes a loader heading, and the list of options and defaults requested through the PARM field of the EXEC statement. The SIZE stated is the size obtained, and not necessarily the size requested in the PARM field. Error messages are written when the errors are detected. After processing is complete, an explanation of the error is written. Loader error messages are similar to those of the linkage editor and are listed in System Messages.

SYSTEM output includes only numbered warning and error messages. These messages are written when the errors are detected. After processing is complete, an explanation of each error is written.

The storage map includes the name and absolute address of each control section and entry point defined in the loaded program. Each map entry marked with an asterisk (*) comes from the data set specified on the SYSLIB DD statement. Two asterisks (***) indicate the entry was found in the link pack area; three asterisks (****) indicate the entry comes from text that was preloaded by a compiler. The TYPE column indicates what each entry on the map is used for: SD=control section, LR=label reference, and PR=pseudoregister.

The map is written as the input to the loader is processed, so all map entries appear in the same sequence in which the input ESD items are defined. The total size and storage extent of the loaded program are also included. For PL/I programs, a list is written showing pseudoregisters with their addresses assigned relative to zero. Figure 76 on page 186 shows an example of a module map. The loader issues an informational message when the loaded program terminates abnormally.

OPTIONS USED-PRINT,MAP,NOLET,CALL,NORES,SIZE=424176

NAME	TYPE	ADDR	NAME	TYPE	ADDR	NAME	TYPE	ADDR	NAME	TYPE	ADDR	NAME	TYPE	ADDR
SAMPL2B	SD	161E0	SAMPL2BA	SD	16EC8	IHEMAIN	SD	17CF8	IHENTRY	SD	17D00	IHESPR	SD	17D10
SYSIN	SD	17D48	IHEVQC	SD	17D80	IHEVQCA	LR	17D80	IHEVQB	SD	17FD8	IHEVQBA	LT	17FD8
IHEDIA	LR	183C0	IHEDIAA	LR	183C0	IHEDIAB	LR	183C2	IHEVPE	SD	18608	IHEVPEA	LR	18608
IHEVPA	LR	18870	IHEVPAA	LR	18870	IHEVFC	SD	189D0	IHEVPCA	LR	189D0	IHEVPC	SD	189F8
IHEVPCA	LR	189F8	IHEVFE	SD	18BE8	IHEVFEA	LR	18BE8	IHEVSC	SD	18C08	IHEVSCA	LR	18C08
IHEDNC	LR	18CB8	IHEDNCA	LR	18CB8	IHEDOA	SD	18F30	IHEDOAA	LR	18F30	IHEDOAB	LR	18F32
IHEDMA	SD	19010	IHEDMAA	LR	19010	IHEVPD	SD	19108	IHEVPDA	LR	19108	IHEVFA	SD	19160
IHEVPAA	LR	19160	IHEVPB	SD	19248	IHEVPBA	LR	19248	IHEXIS	SD	193F0	IHEXISO	LR	193F0
IHEIOB	SD	19488	IHEIOBA	LR	19488	IHEIOBB	LR	19490	IHEIOBC	LR	19498	IHEIOBD	LR	194A0
IHESARC	LR	1A9CB	IHESADD	LR	1A9DE	IHESAFF	LR	1AA18	IHEPRT	SD	1AB70	IHEPRTA	LR	1AB70
IHEBEGA	LR	1AE28	IHEERR	SD	1AE68	IHEERRD	LR	1AE68	IHEERRC	LR	1AE72	IHEERRB	LR	1AE7C
IHEERRA	LR	1AE86	IHEERRE	LR	1B4E2	IHEIOF	SD	1B580	IHEIOFB	LR	1B580	IHEIOFA	LR	1B582
IHEITAZ	LR	1B81E	IHEITAX	LR	1B82A	IHEITAA	LR	1B83E	IHEDCN	SD	1B860	IHEDCNA	LR	1B860
IHEDCNB	LR	1B862	IHEIOD	SD	1BA50	IHEIODG	LR	1BA50	IHEIODP	LR	1BA52	IHEIODT	LR	1BB4A
IHEVTB	SD	1BCF0	IHEVTBA	LR	1BCF0	IHEVQA	SD	1BD78	IHEVQAA	LR	1BD78			
IHEQINV	PR	00	IHEGERR	PR	4	SAMPL2BB	PR	8	SAMPL2BC	PR	C	IHEQSPR	PR	10
SYSIN	PR	14	IHEQLSA	PR	18	IHEQLWO	PR	1C	IHEQLW1	PR	20	IHEQLW2	PR	24
IHEQLW3	PR	28	IHEQLW4	PR	2C	IHEQLWE	PR	30	IHEQLCA	PR	34	IHEQVDA	PR	38
IHEQFVD	PR	3C	IHEQCFL	PR	40	IHEQPOP	PR	48	IHEQADC	PR	4C	IHEQXLV	PR	50
IHEQEVT	PR	58	IHEQSLA	PR	60	IHEQSAR	PR	64	IHEQLWF	PR	68	IHEQRTC	PR	6C
IHEQSFC	PR	70												
IEW1001	IHEUPBA													
IEW1001	IHEUPAA													
IEW1001	IHETERA													
IEW1001	IHEM91C													
IEW1001	IHEM91B													
IEW1001	IHEM91A													
IEW1001	IHEDDOD													
IEW1001	IHEVPPA													
IEW1001	IHEVPDA													
IEW1001	IHEDBNA													
IEW1001	IHEVSFA													
IEW1001	IHEVSBA													
IEW1001	IHEVCAA													
IEW1001	IHEVSAA													
IEW1001	IHEDNBA													
IEW1001	IHEUPBB													
IEW1001	IHEUPAB													
IEW1001	IHEVSEB													

TOTAL LENGTH 5068
ENTRY ADDRESS 17D00

IEW1001 WARNING - UNRESOLVED EXTERNAL REFERENCE (NOCALL SPECIFIED)

Figure 76. Module Map Format Example

APPENDIX D. LOADER STORAGE CONSIDERATIONS

The loader requires virtual storage space for the following items:

- Loader code
- Data management access methods
- Buffers and tables used by the loader (dynamic storage)
- Loaded program (dynamic storage)

Region size includes all four of the above items; the SIZE option refers to the last two items.

For the SIZE option, the minimum required virtual storage is 4K bytes plus the size of the loaded program. This minimum requirement grows to accommodate the extra table entries needed by the program being loaded. For example, FORTRAN requires at least 3K bytes plus 4K bytes plus the size of the loaded program, and PL/I needs at least 8K bytes plus 4K bytes plus the size of the loaded program. Buffer number (BUFNO) and block size (BLKSIZE) could also increase this minimum size. Figure 77 shows the appropriate storage requirements in bytes.

The maximum virtual storage that can be used is whatever virtual storage is available up to 8192K bytes.

All or part of the storage required is obtained from user storage. If the access methods are made resident at IPL time, they are allocated in system storage. However, 6K bytes is always reserved for system use.

The loader code could also be made resident in the link pack area. If so, it requires the following space: HEWLDRGO, the control/interface module (alias LOADER), approximately 700 bytes; HEWLOADR, the loader processing portion, approximately 13 664 bytes.

The size of the loaded program is the same as if the program had been processed by the linkage editor and program fetch.

The loader does not use auxiliary storage space for work areas.

Consideration	Approximate Virtual Storage Requirements (in Bytes)	Comments
Loader Code Control	2000	
Loader Code Processing	14000	
Data Management	6K	BSAM
Object Module Buffers and DECBS	$BUFNO \times (BLKSIZE + 24)$	Concatenation of different BLKSIZE and BUFNO must be considered. (Minimum BUFNO=2)
Load Module Buffer and DECBS	304	

Figure 77 (Part 1 of 2). Virtual Storage Requirements

Consideration	Approximate Virtual Storage Requirements (in Bytes)	Comments
SYSTEM DCB Buffers and DECBs	312	Allocated if TERM option is specified
SYSLOUT Buffers and DECBs	BUFNO*(BLKSIZE + 24)	Buffer size rounded up to integral number of double words. (Minimum BUFNO=2)
Size of program being loaded	Program size	Program size is restricted only by available virtual storage, up to 8 megabytes
Each external relocation dictionary entry	8	
Each external symbol	20	
Largest ESD number	4n (n is the largest number of ESDs in any input module)	Allocated in increments of 32 entries
Fixed Loader Table Size	1260	Subtract 88 if NOPRINT is specified
Condensed Symbol Table	12n (n is the total number of control sections and common areas in the loaded program)	Built only if TSO is operating and space is available
System Requirements	4000	

Figure 77 (Part 2 of 2). Virtual Storage Requirements

APPENDIX E. LOADER RETURN CODES

The return code of a loader step is determined by the return codes resulting from loader processing and from loaded program processing.

The return code indicates whether errors occurred during the execution of the loader or of the loaded program. The return code can be tested through the COND parameter of the JOB statement specified for this job and/or the COND parameter of the EXEC statement specified in any succeeding job step. (For details, see the publication JCL). Figure 78 shows the return codes.²

Code Returned to Caller	Loader Return Code	Program Return Code	Conclusion or Meaning
0	0	0	Program loaded successfully, and execution of the loaded program was successful.
0	4	0	The loader found a condition that may cause an error during execution, but no error occurred during execution of the loaded program.
0	8(LET)	4	The loader found a condition that may cause an error during execution, but no error occurred during execution of the loaded program.
4	0	4	Program loaded successfully, and an error occurred during execution of the loaded program.
4	4	4	The loader found a condition that may cause an error during execution, and an error did occur during execution of the loaded program.
4	8(LET)	4	The loader found a condition that may cause an error during execution, and an error did occur during execution of the loaded program.
8	0	8	Program loaded successfully, and an error occurred during execution of the loaded program.

Figure 78 (Part 1 of 2). Return Codes

² Error diagnostics (SYSOUT and/or SYSTEM data set) for the loader will show the severity of errors found by the loader.

Code Returned to Caller	Loader Return Code	Program Return Code	Conclusion or Meaning
8	4	8	The loader found a condition that may cause an error during execution, and an error did occur during execution of the loaded program.
8	8(LET)	8	The loader found a condition that may cause an error during execution, and an error did occur during execution of the loaded program.
8	8		The loader found a condition that could make execution impossible. The loaded program was not executed.
12	0	12	Program loaded successfully, and an error occurred during execution of the loaded program.
12	4	12	The loader found a condition that may cause an error during execution, and an error did occur during execution of the loaded program.
12	8(LET)	12	The loader found a condition that may cause an error during execution, and an error did occur during execution of the loaded program.
12	12		The loader could not load the program successfully; execution impossible.
16	0	16	Program loaded successfully, and the loaded program found a terminating error.
16	4	16	The loader found a condition that may cause an error during execution, and a terminating error was found during execution of the loaded program.
16	8(LET)	16	The loader found a condition that may cause an error during execution, and a terminating error was found during execution of the loaded program.
16	16		The loader could not load program; execution impossible.

Figure 78 (Part 2 of 2). Return Codes

GLOSSARY

This glossary includes definitions developed by the American National Standards Institute (ANSI). This material is reproduced from the American National Dictionary for Information Processing, copyright 1977 by the Computer and Business Equipment Manufacturers Association, copies of which may be purchased from the American National Standards Institute, 1430 Broadway, New York, New York 10018. ANSI definitions are preceded by an asterisk(*).

***address.** An identification, as represented by a name, label, or number, for a register, location in storage, or any other data source or destination such as the location of a station in a communication network; any part of an instruction that specifies the location of an operand for the instruction.

address constant. A value, or an expression representing a value, used in the calculation of storage addresses; can be used for branching or retrieving data.

addressing mode (AMODE). The attribute of an entry point in which control is received.

address translation. The process of changing the address of a data item or an instruction from its virtual address to the real storage address of the location where it will reside. See also dynamic address translation.

alias name. An alternate name or entry point for a load module that is also entered in the output module library directory entry along with the member name.

automatic call library mechanism. The process whereby control sections are processed by the linkage editor or loader to resolve external references to members of partitioned data sets not resolved by primary input processing.

auxiliary storage. Data storage other than virtual storage; for example, storage on magnetic tape or direct-access devices.

common area. A control section used to reserve a virtual storage area that can be referred to by other modules; may be either named or unnamed (blank).

common segment. A segment upon which two exclusive segments are dependent.

composite external symbol dictionary (CESD). Control information associated

with a load module that identifies the external symbols in the module.

control section. That part of a program (instructions and data) specified by the programmer to be a relocatable unit, all elements of which are to be loaded into adjoining storage locations for execution. Abbreviated CSECT.

control section name. The symbolic name of a control section.

demand paging. Transfer of a page from external page storage to real storage at the time it is needed for execution.

downward reference. A reference made from a segment to another segment lower in the same path; that is, farther from the root segment.

dynamic address translation (DAT). (1) The change of a virtual storage address to a real storage address during execution of an instruction. See also address translation. (2) A hardware feature that performs the translation.

entry name. A name within a control section that defines an entry point, and can be referred to for execution by any control section.

exclusive reference. A reference between exclusive segments; that is, a reference from a segment in storage to an external symbol in a segment that will cause overlay of the calling segment.

exclusive segments. Segments in the same region of an overlay program, neither of which is in the path of the other; they cannot be in virtual storage simultaneously.

external name. A name that can be referred to by any control section or separately assembled or compiled module; that is, a control section name or an entry name.

external page storage. The portion of auxiliary storage that is used to contain pages.

external reference. (1) A reference to a symbol that is defined as an external name in another module. (2) An external symbol that is defined in another module; that which is defined in the Assembler language by an EXTRN statement or by a V-type address constant, and is resolved during linkage editing. See also weak external reference.

external symbol. A control section name, entry point name, or external reference that is defined or referred to in a particular module. A symbol contained in the external symbol dictionary.

external symbol dictionary (ESD). Control information associated with an object or load module that identifies the external symbols in the module.

inclusive reference. A reference between inclusive segments; that is, a reference from a segment in storage to an external symbol in a segment that will not cause overlay of the calling segment.

inclusive segments. Segments in the same region of an overlay program that are in the same path; they can be in virtual storage simultaneously.

invalid exclusive reference. An exclusive reference in which a common segment does not contain a reference to the symbol used in the exclusive reference.

library. In this publication, a partitioned data set that always contains named members.

load module. The output of the linkage editor; a program in a format suitable for loading into virtual storage for execution.

load module buffer. An entity of virtual storage used by the linkage editor to read input load module text records and possibly to retain the text information in storage for subsequent writing of the output load module text records.

***module.** A program unit that is discrete and identifiable with respect to compiling, combining with other units, and loading, for example, the input to, or output from, an assembler, compiler, linkage editor, or executive routine.

multiple load module processing. A method of processing whereby two or more load modules can be produced in a single linkage editor job step.

***object module.** A module that is the output of an assembler or compiler and is input to a linkage editor.

overlay program. A program in which certain control sections can use the same storage locations at different times during execution.

***overlay supervisor.** A routine that controls the proper sequencing and positioning of segments of computer programs in limited storage during their execution.

overlay tree. A graphic representation showing the relationships of segments of an overlay program and how the segments are arranged to use the same main storage area at different times.

page. (1) A fixed-length block of instructions, data, or both, that can be transferred between real storage and external page storage. (2) To transfer instructions, data, or both between real storage and external page storage.

page fault. A program interruption that occurs when a page that is marked "not in real storage" is referred to by an active page.

paging. The process of transferring pages between real storage and external page storage.

path. All of the segments in an overlay tree between a given segment and the root segment, inclusive.

private code. An unnamed control section.

program. A logically self-contained sequence of operations or instructions that, when followed in some predetermined sequence, will produce a specified result; a sequence of instructions to be performed by a computer; one or more modules, in source language or relocatable object code, or one module in executable code, that is a logically self-contained process.

program fetch. A program that prepares load modules for execution by loading them at specific storage locations; it also readjusts each address constant.

pseudoregister. In PL/I, a location in virtual storage that is used as a pointer to dynamically acquired virtual storage. It enables the PL/I compiler to generate reenterable code. External dummy sections give the programmer using Assembler F or Assembler H the same facility.

real storage. The storage from which the central processing unit can directly obtain instructions and data, and to which it can directly return results.

reenterable load module. A module that can be used concurrently by more than one task.

refreshable load module. A load module that cannot be modified by itself or by any other module during execution; can be replaced by a new copy during execution by a recovery management routine without changing either the sequence or results of processing.

region. In an overlay structure, a contiguous area of virtual storage

within which segments can be loaded independently of paths in other regions. Only one path within a region can be in virtual storage at any one time.

relocation. The modification of address constants required to compensate for a change of origin of a module, program, or control section.

RSECT. A read-only CSECT in the nucleus.

root segment. That segment of an overlay program that remains in virtual storage at all times during the execution of the overlay program; the first segment in an overlay program.

residence mode (RMODE). Defines whether the program must be resident in storage addressable by 24-bit addressing or 31-bit addressing.

scatter format. A load module attribute that permits the control program to dynamically load control sections into noncontiguous areas of virtual storage.

segment. The smallest functional unit (one or more control sections) that can be loaded as one logical entity during execution of an overlay program.

serially reusable load module. A module that cannot be used by a second task until the first task has finished using it.

source module. The source statements that constitute the input to a language translator for a particular translation.

storage block. A 2K-byte block of real storage to which a storage key can be assigned, processor-model dependent.

upward reference. A reference made from a segment to another segment higher in the same path; that is, closer to the root segment.

valid exclusive reference. An exclusive reference in which a common segment contains a reference to the symbol used in the exclusive reference.

virtual address. An address that refers to virtual storage and must, therefore, be translated into a real storage address when it is used.

virtual storage. Addressable space that appears to the user as real storage, from which instructions and data are mapped into real storage locations. The size of virtual storage is limited by the addressing scheme of the computing system and the amount of auxiliary storage available, rather than by the actual number of real storage locations.

weak external reference. An external reference that does not have to be resolved during linkage editing. If it is not resolved, it appears as though its value was resolved to zero. Abbreviated WXTRN.

INDEX

Special Characters

\$PRIVATE 119
**GO 172

A

A-type address constant
 replacing control sections 90, 98
 SEGWT macro 164
AC option 18, 41
adcons
 See address constant
additional call libraries 31
additional input sources
 automatic call library 28-32
 general description of 14, 23
 included data sets 33-36
 libraries 28-36
 processing of 28-36
 specification of
 automatic call library 30
 INCLUDE statement 34-36
 LIBRARY statement 30, 79
address
 assignment 10
 of main entry point
 module map 119
 specifications 112
 sequence in object module text 8
address constant 6, 9
 See also A-type, V-type address
 constant
 resolution of 6-9
addressing mode
 assignment
 linkage editor 18
 combinations
 residence mode 43
 control section name 7
 default 18
 entry point 113
 linkage editor 1
 options 157
 override 18
 parameter
 linkage editor 42
 private code 7
alias name
 example 111
 linkage editor 37
 loader 180
 specifying 111
ALIAS statement 69, 111
alignment, page 105
alternate output data set
 See SYSTERM data set
AMODE
 See addressing mode
APF (Authorized Program Facility) 17
assigning block size, linkage editor 58
asynchronous overlay supervisor 160
attributes, module

 See module attributes
authorization codes
 See AC option
Authorized Program Facility
 See APF
automatic call library for loader
 DD statement for 174
 description of 166
 options 171
automatic call library mechanism 119
 See also call library, linkage editor
 module map 119
automatic deletion of modules 166
automatic replacement
 control sections 98-100
 examples 98-100
 modules 111
 note on overlay programs 98
automatic search
 of link pack area queue 173
 of SYSLIB 171

B

blank common area
 collection of 113, 157
 definition 7
 module map 118
BLKSIZE operand (DCB macro) 55-61
block size 55-61
blocking factors, SIZE option 49, 61
branch instructions, in overlay
 programs 160-162
buffer, load module
 See load module buffer
BUFNO operand (DCB macro)
 loader data sets 174
 loader DD statements 174

C

call library, linkage editor
 additional libraries 31
 concatenating 30
 ddname 29
 example 29
 NCAL option 32
 negating 32, 45
 never-call 32
 restricted no-call 31
 specification of 29-30
call library, loader
 DD statement for 174-176
 description 166
 options for use 171
CALL loader option 171
CALL macro
 definition 161
 invoking the loader 182
 with only loadable attribute 39
capacities, linkage editor 136
cataloged procedure

- adding DD statements 66
- definition 62
- linkage editor 62
- LKED 62-64
- LKEDG 64-65
- overriding 65-66
- CESD (composite external symbol dictionary)
 - definition 10
 - number of entries permitted 137
- CHANGE statement
 - example 96
 - summary 70-71
 - using INCLUDE statement 103
 - using REPLACE statement 103
- changing external symbols 96
- class test table 147
- collection of common areas 113, 157-159
- common area
 - blank 7
 - collection of 113, 157, 159
 - definition 7
 - lengthening 17, 73
 - module map 118
 - named 7
 - ordering named 103-105
 - reserving storage for 113
- common segment
 - definition 145
 - exclusive references 146
 - promotion of common areas 157
- comparison of linkage editor and loader 1, 166
- compatibility, linkage editor and loader 169
- composite external symbol dictionary
 - See CESD
- concatenation
 - call libraries 30
 - data sets
 - linkage editor 35
 - loader 175
 - restrictions 60
- COND parameter
 - determining load module execution 54
 - in LKEDG 64
 - specified in EXEC statement 54
 - specified in JOB statement 54
- condition parameter
 - See COND parameter
- constant
 - See address constant
- control dictionaries 6
- control section
 - aligning on page boundary 105
 - definition 5
 - deleting 102
 - external symbol dictionary 6
 - lengthening 17, 73
 - module map 118
 - name
 - changing 96
 - external symbol dictionary 6
 - ordering of 103-105
 - positioning 153
 - replacing 97
- control statements
 - as input 26
 - concatenating object module data set 26
 - continuation of 67
 - format conventions 67
 - general format 67
 - listing 118

- listing option 51
- placement information 68
- summary list 67
- cross-reference table
 - option 52
 - producing 119
 - sample 120
- CSECT identification records
 - function 17
 - object and load modules 6
 - storage required 137
 - using IDENTIFY statement 74

D

- data definition statement
 - See DD statement
- data for loaded program 177
- data set
 - concatenated 30, 35, 175-176
 - linkage editor
 - input 23
 - output 109
 - loader 175
- DCB information
 - linkage editor 55
 - loader 174
- DCBS option, block size 50
- DD statement
 - general description 55
 - linkage editor data sets
 - ddnames 55-57
 - SYSLIB 31, 57
 - SYSLIN 56
 - SYSLMOD 58
 - SYSPRINT 57
 - SYSTEM 59
 - SYSUT1 57
 - loader data sets
 - ddnames 174, 180
 - SYSLIB 175
 - SYSLIN 175
 - SYSLOUT 176
 - SYSTEM 176
- ddname list 108
- ddnames, linkage editor
 - invoking 55-61
 - loader
 - automatic call library 174
 - diagnostic data set 176
 - input data set 175
 - specifying alternate names 108, 180
- default module attributes 43
- deleting a control section 102
- deleting an entry name 102
- deleting modules 166
- diagnostic messages
 - linkage editor
 - directory 115
 - format 115-117
 - loader
 - defined by SYSLOUT DD and SYSTEM DD 185
 - format 185
- diagnostic output
 - linkage editor
 - messages 115
 - optional 117-120
 - options, summary 16-17
 - loader

- data set 176
 - format 185
 - options 171
- dictionaries
 - composite external symbol 10, 137
 - external symbol 7
 - relocation 6, 9, 136
- directory entry
 - authorization code 113
 - changing 97
- disposition messages 115-116
- downward call
 - See downward reference
- downward reference
 - definition 139
 - maximum number 138

E

- editing
 - conventions 94-95
 - module 94-95
- end of module indication 9
- END statement
 - object module 6
 - specifying entry point 112
- ENTAB (entry table) 148
- entry address, module map 119
- entry name
 - definition 7
- ESD
 - changing 96
 - deleting 102
 - module map 118
- entry point
 - example 113
 - loaded program 171, 182
 - specification of
 - END statement 112
 - ENTRY statement 72, 112
- ENTRY statement
 - main entry point 112
 - summary 72
- entry table 148
- EOM (end of module indication) 9
- EP loader option 171
- error conditions
 - See severity code
- error messages
 - See diagnostic messages
- ESD (external symbol dictionary) 7
- exclusive call option 44
- exclusive reference
 - definition 145
 - entry table 148
 - restrictions 146
 - segment table 148
 - XCAL option 44
- exclusive segments
 - note on overlay programs 98
- EXEC statement
 - linkage editor
 - introduction 37
 - job step options 37
 - program name 37
 - REGION parameter 54
 - return code 54
 - loader
 - description 170, 174
 - examples 173
- executable module 44

- EXPAND statement 17, 73
- external dummy section 7, 15, 114
 - See also pseudoregister
 - definition 7
 - processing of 15, 114
- external name 5
 - See also control section name, entry name
 - definition 5
- external reference
 - changing 96
 - definition 5, 7
 - external symbol dictionary 7
 - resolving 11, 28
 - weak
 - automatic library-call 28
 - cross-reference table 119
- external symbol
 - changing 96
 - definition 5
- external symbol dictionary 7

F

- functions
 - linkage editor 12
 - loader 166

H

- HEWL 37, 62
- HEWLOAD 181, 184
- HEWLOADR 181

I

- IDENTIFY statement summary 74
- IDR
 - See CSECT identification records
- IEBUPDTE program
 - input statements 132
- INCLUDE statement
 - requesting additional data sets as
 - input 33
 - specifying ddname of DD statement 33
 - summary 76
- included data sets
 - concatenated data sets 34
 - library members 35
 - linkage editor 33
 - sequential data sets 34
- inclusive reference
 - when to use 145
- inclusive segments
 - definition 145
- incompatible job step options 52
- incompatible module attributes 44
- input data sets
 - linkage editor 23
 - loader 175
- input processing 23
- input sources
 - linkage editor 9
 - loader 170, 175
- INSERT statement

- summary 77
- using 155
- intermediate data set
 - devices supported 138
 - linkage editor 9
 - loader 169
 - SIZE option 45, 136
 - storage required 136
 - SYSUTI DD statement 57
- intermediate text records, number produced 137
- internal data area 169
- invalid attributes or options 115
- invalid exclusive reference illustration 146
- invocation of
 - linkage editor 107
 - loader 180

J

- job control language summary 37
- job control statements
 - linkage editor 37
 - loader processing
 - basic format 170
 - compile-load job 178
 - load job 189
 - multiple compilations 178
- job step
 - options
 - EXEC statement 37

L

- let execute option 44
- LET option
 - linkage editor 44, 129
 - loader 169, 172
 - overlay programs 157
- library call
 - See automatic call library for loader
- library members
 - including 35
 - input to linkage editor 24
 - input to loader 175
- LIBRARY statement
 - additional call libraries 31
 - NCAL option 45
 - never-call function 32
 - restricted no-call function 31
 - summary 79
 - using 30
- LINK command 22
- LINK macro
 - invoking
 - linkage editor 107
 - loader 181
- link pack area resolution loader 173
- linkage editor
 - assigning block size 58
 - capacities 136
 - cataloged procedures 62
 - compared to loader 1, 166
 - control statement summary 67
 - DD statements 56-61

- functions 12
- input
 - additional data sets 23
 - control statements 27
 - object modules 27
 - primary data sets 23
- invoking 107
- output 109
- processing 9
- relationship to operating system 21
- storage requirements 136
- when to use 1
- LINKEDIT 37
- linking modules 13
- LIST option 51, 118
- LKED procedure 62-64
- LKEDG procedure 64-65
- LOAD macro
 - invoking the loader 183
 - only loadable modules 39
- load module
 - attribute assignment 16
 - attributes 38
 - buffer 46
 - definition 4
 - entry point 112
 - input
 - linkage editor 23
 - loader 170
 - linkage editor output 109
 - multiple processing of 114
 - structure 6
- load module buffer
 - allocating storage 46
- load module creation 10
- load module format loader
 - example 121
- load point 144, 151
- load step 1, 170
- loaded program
 - data 177
 - module map 185
 - options 171, 174
 - return codes 189
- loader
 - abnormal termination message (VS2) 185
 - alias name 180
 - compared to linkage editor 1, 166
 - compatibility with linkage editor 169
 - data sets 175
 - input 166, 170
 - invoking 180
 - options 171, 174
 - output 185
 - program name 180
 - restrictions 169
 - return codes 189
 - when to use 1
- LOADGO command 169
- loading
 - with identification 184
 - without identification 182
- logical record length
 - linkage editor data sets
 - blocking factors 56
 - diagnostic output 58
 - input 56
 - SIZE option 45
- LRECL operand (DCB macro) 55-56

M

macros, linkage editor
 format 107
 MAP option
 linkage editor 51, 118-119
 loader 169, 172
 maximum record size for device
 types 46-47
 member name
 definition 110
 example 110
 specifying 110
 member, partitioned data set
 including 35
 input to the linkage editor 24
 input to the loader 175
 messages
 disposition 115-116
 examples 117
 format 116
 text 117
 unnumbered 116
 MODE statement
 example 82
 specifying addressing mode 81
 specifying residence mode 81
 summary 81-82
 values 81
 modular programming 4
 module attributes
 default attributes 43
 describing output module 38
 downward compatible 38
 incompatible attributes 44, 52
 not editable 39
 not-executable 43
 only loadable 39
 overlay 39
 refreshable 41
 reusability
 reenterable 40
 serially reusable 40
 scatter format 38
 test 41
 module disposition messages 115
 module editing
 general description 94
 summary 13
 module linking 13-14
 module map
 linkage editor
 description 118-119
 example 119
 MAP option 51
 loader
 description 185
 example 186
 specifying 172
 module, definition 4
 multiple load module processing
 producing 114
 multiple region overlay program
 general description 148
 specifying 152
 using 148

N

NAME loader option 172
 NAME statement
 multiple load module processing 114
 replace function 111
 summary 83
 SYSLMOD DD 110
 named common area
 aligning on page boundary 105
 collection of 113, 157
 definition 7
 module map 118
 NCAL option
 linkage editor 32, 45
 loader 169, 171
 NE attribute 39
 negation of automatic library call
 linkage editor 32
 loader
 diagnostic output 173
 module map 173
 search of link pack area 173
 not editable attribute 39
 not executable attribute 45
 reenterable attribute 40
 refreshable attribute 41
 serially reusable 40
 never-call function
 cross-reference table 119
 specifying external references 32
 no automatic library call option 45
 no-call 32
 NOCALL loader option 171
 node point
 See load point
 NOLET loader option 172
 NOMAP loader option 172
 NOPRINT loader option 172
 NORES loader option 173
 not editable attribute
 linkage editor 39
 loader 169
 not-executable attribute 44
 NOTERM loader option 173

O

object module
 definition 4
 input to linkage editor 27
 input to loader 175
 structure 6
 virtual storage 169
 OL attribute 39
 only loadable attribute 39
 optional output 117
 options, incompatible 52
 options, linkage editor
 addressing mode 157
 module attributes 38
 output 51
 residence mode 157
 space allocation 45
 special processing 44-45
 ORDER statement 84, 103-105
 origin
 control section in module map 118
 region 152

- segments 144
- output module library 109
- output of linkage editor
 - diagnostic messages 115
 - load module 109
 - optional output 117-120
 - output module library 109
 - output options 51
- output of the loader
 - messages 185
 - module map 185
 - specifying 170-174
- output text record length 136
- overlap of loading and processing,
 - overlay segments 162
- overlay attribute
 - specifying 39
- overlay program
 - communication 160
 - design 139
 - module map 118
 - multiple region 148
 - process 147-148
 - region origin 152
 - respecifying control statements 151
 - sample program 127-132
 - segment origin 144, 151
 - single region 140
 - special considerations 157
 - specifying 150
 - storage requirements 159-160
- OVERLAY statement
 - specifying 150-156
 - summary 86
- overlay supervisor, definition 147
- overlay tree
 - positioning segments 142
- overriding cataloged procedures
 - DD statements 66
 - EXEC statement 65
- OVLY attribute 39-40

P

- page boundary
 - aligning control sections or named
 - common areas 105
- PAGE statement
 - aligning control sections 105
 - summary 88
- parameter
 - addressing mode 42
 - combination 42
 - residence mode 42
- partitioned data set
 - input
 - linkage editor 24
 - loader 175
 - output, linkage editor 109
- path
 - in overlay program 139
- placement of control statements 68
- positioning control sections 153
- preloaded text 169, 185
- primary input data set
 - control statements 26
 - job control language
 - specifications 23
 - object modules 23, 27
- PRINT loader option 172
- private call libraries 30

- private code
 - definition 7
 - module map 118
- procedure LKED 62-64
- procedure LKEDG 64-65
- processing history, tracing
 - CSECT identification record 17
- processing options, special 44
- program fetch
 - functions 11
- prompter
 - linkage editor, function of 22
 - loader, function of 169
- pseudoregister
 - definition 7
 - module map 119
 - processing of 15, 114

R

- read-only attribute, assignment 21
- RECFM (record format)
 - linkage editor data sets
 - diagnostic output 61
 - input 55-56
 - load modules 56-61
- record format
 - See RECFM
- record size
 - maximum
 - device type 46
- reenterable attribute 40
- reenterable load module
 - module attribute 40
- REFR attribute 41
- refreshable attribute 41
- refreshable load module
 - module attribute 44
- REGION parameter
 - specifying storage 54
- region, overlay programs
 - specifying 152
 - using 148
- region, virtual storage
 - linkage editor
 - cataloged procedures 62
 - SIZE option 50
 - loader 182
- relocating a load module 4
- relocation dictionary
 - See RLD
- RENT attribute 40
- replace function 97-103, 111
- REPLACE statement
 - deleting CSECT 103
 - example 101
 - sample program 123-127
 - summary 90-91
 - using 101
- replacing control sections, assembler
 - language note 98
- replacing external symbols
 - See CHANGE statement, changing
 - external symbols
- replacing load modules with the same
 - name 111
- repositioning control statements
 - automatic call library 156
 - INSERT control statement 77, 153
- reprocessing load modules
 - entry point assignment 112

- not editable attribute 39
- RES loader option 173
- reserving storage 113
- residence mode
 - assignment
 - linkage editor 19
 - output load module 109
 - combinations
 - addressing mode 43
 - control section name 7
 - default 19
 - entry point 113
 - options 157
 - override 19
 - parameter
 - linkage editor 43
 - private code 7
- resolving external references 11, 28
- restricted no-call function 31
- restrictions, virtual storage size requirements 41
- return codes
 - linkage editor 54
 - loader 189
 - severity code 116
- REUS attribute 40
- reusability attributes
 - general description 40
 - reenterable 40
 - serially reusable 40
- RLD (relocation dictionary)
 - number of entries 136
 - using 9
- RMODE
 - See residence mode
- root segments
 - definition 139
 - OVERLAY 151
 - segment table 148

S

- sample programs 122
- scatter loading 38
- SCTR attribute 38
- SEGLD macro 160
- segment 142, 144, 145, 147
 - See also exclusive, inclusive, root segments
 - communication 145-147
 - dependency 142
 - origin 144
- segment load macro 162-163
- segment table 148
- segment wait macro
 - SEGLD 163
 - using 163
- SEGTAB (segment table) 148
- SEGWT macro
 - SEGLD 163
 - using 163
- sequential data set
 - INCLUDE statement 34
 - input to linkage editor 23, 34
 - input to loader 175
- serially reusable
 - attribute 40
- SETCODE statement 18, 92
- SETSSI statement 93
- severity code
 - linkage editor messages 116

- return codes 54
- severity 0,2 errors 116
- SIZE option
 - linkage editor 45, 61
 - loader 173, 187
- space allocation
 - DCBS option 50
 - maximum values 45, 48
 - minimum values 45, 48
 - SIZE option 45
 - specifying storage 45
- special processing options
 - affecting automatic library call mechanism 44
 - affecting output module 44
 - summary 16
- static external areas 113
- SYSLIB DD statement
 - automatic call library 28
 - linkage editor 57
 - loader 175
- SYSLIN DD statement 56, 175
 - See also automatic call library for loader
 - linkage editor 56, 175
- SYSLMOD DD statement 58, 114, 115
 - See also output module library function 58
 - NAME statement 114-115
 - referenced by INCLUDE statement 58
- SYSLOUT DD statement 172, 176
- SYSPRINT DD statement 57
- system call library
 - example 29
 - list 29
- system status index information, storage of 17
- SYSTEM data set
 - linkage editor 52, 59, 117
 - loader 174, 176, 185
- SYSTEM DD statement
 - linkage editor 52, 59, 117
 - loader 174, 176, 185
- SYSUT1 DD statement 57

T

- TEMPNAME 110
- temporary data set 25, 110
- TERM option
 - linkage editor 52, 60, 117
 - loader 173
- TEST attribute 41
- text
 - message 117
 - note 8
 - processing out-of-order 6
- time sharing option
 - See TSO
- tracing processing history 17
- TRANSFORM table 147
- tree structure
 - positioning of segments 142
 - purpose of 141
- TSO (time sharing option)
 - linkage editor
 - invoking 22
 - SYSTEM data set 59
 - TERM option 117
 - loader
 - invoking 169

SYSTEM data set 173, 176
TERM option 173

U

unnumbered messages 115-116
unresolved references
 automatic library-call, resolving
 with 28
 cross-reference table 119
upward reference, definition 139
user-specified
 input 9
 storage 16
user-written library 30

V

V-type address constant
 branch instruction, overlay 162
 CALL 162
 SEGLD 163
 SEGWT 164
valid exclusive reference 146

virtual storage requirements
 linkage editor 136
 loader 187
 overlay programs 159-160
 restrictions 41

W

wait for loading of segment 163
warning messages 116-117
weak external reference
 automatic library-call 28
 cross-reference table 119
 definition 7
 level F linkage editor 14

X

XCAL option 44, 157
XCTL macro
 input to loader 169
 invoking the loader 180
XREF option 52



This manual is part of a library that serves as a reference source for system analysts, programmers, and operators of IBM systems. You may use this form to communicate your comments about this publication, its organization, or subject matter, with the understanding that IBM may use or distribute whatever information you supply in any way it believes appropriate without incurring any obligation to you.

Your comments will be sent to the author's department for whatever review and action, if any, are deemed appropriate.

Note: Do not use this form to request IBM publications. If you do, your order will be delayed because publications are not stocked at the address printed on the reverse side. Instead, you should direct any requests for copies of publications, or for assistance in using your IBM system, to your IBM representative or to the IBM branch office serving your locality.

Note: Staples can cause problems with automated mail sorting equipment.
Please use pressure sensitive or other gummed tape to seal this form.

If you have applied any technical newsletters (TNLs) to this book, please list them here:

Last TNL _____

Previous TNL _____


Fold on two lines, tape, and mail. No postage stamp necessary if mailed in the U.S.A.
(Elsewhere, an IBM office or representative will be happy to forward your comments or you may mail directly to the address in the Edition Notice on the back of the title page.) Thank you for your cooperation.

Reader's Comment Form

Fold and tape

Please do not staple

Fold and tape



NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES

BUSINESS REPLY MAIL
FIRST CLASS PERMIT NO. 40 ARMONK, N.Y.

POSTAGE WILL BE PAID BY ADDRESSEE

IBM Corporation
P.O. Box 50020
Programming Publishing
San Jose, California 95150



Fold and tape

Please do not staple

Fold and tape






MVS/370
Linkage Editor and Loader
User's Guide

File Number S370-40

GC26-4061-01



Printed in U.S.A.